NAVAL POSTGRADUATE SCHOOL
Monterey, California

THESIS

SHIP-SHORE PACKET SWITCHED
COMMUNICATIONS SYSTEM

by

Rex A. Buddenberg

June 1986

Thesis Advisor: Paul Cooper

Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3 DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE	
4 PERFORMING ORGANIZATION REPORT NUMBER(S)	5 MONITORING ORGANIZATION REPORT NUMBER(S)
6a. NAME OF PERFORMING ORGANIZATION Naval Postgraduate School 6b OFFICE SYMBOL (*If applicable*) 54	7a. NAME OF MONITORING ORGANIZATION Naval Postgraduate School
6c. ADDRESS (*City, State, and ZIP Code*) Monterey, CA 93943	7b. ADDRESS (*City, State, and ZIP Code*) Monterey, CA 93943
8a NAME OF FUNDING / SPONSORING ORGANIZATION 8b. OFFICE SYMBOL (*If applicable*)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
8c. ADDRESS (*City, State, and ZIP Code*)	10 SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO / PROJECT NO / TASK NO / WORK UNIT ACCESSION NO

11 TITLE (*Include Security Classification*)
SHIP-SHORE PACKET SWITCHED COMMUNICATIONS SYSTEM

12 PERSONAL AUTHOR(S)
Buddenberg, Rex A.

13a TYPE OF REPORT	13b TIME COVERED	14 DATE OF REPORT (*Year, Month, Day*)	15 PAGE COUNT
Master's Thesis	FROM _____ TO _____	June 1986	230

16 SUPPLEMENTARY NOTATION

17 COSATI CODES			18 SUBJECT TERMS (*Continue on reverse if necessary and identify by block number*)
FIELD	GROUP	SUB-GROUP	Communications System Engineering, Ship-shore, HF Communications, Packet Switching

19 ABSTRACT (*Continue on reverse if necessary and identify by block number*)

This thesis presents an architecture for ship-shore sea service communications. Starting with the ARPANET packet switching model (TCP/IP), Network and Logical Link layers are defined which deal with the following problems:

1) Many ship-shore links are one way only. A Network level acknowledgement protocol to work under TCP/IP, integrates multiple, heterogenous, one-way links into a complete network.

2) Conventional network access procedures are invalid in HF ship-shore communications because the assumption that ships can hear each other transmit cannot be made. Two network access techniques are presented.

20 DISTRIBUTION / AVAILABILITY OF ABSTRACT	21. ABSTRACT SECURITY CLASSIFICATION
☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS	UNCLASSIFIED
22a NAME OF RESPONSIBLE INDIVIDUAL Professor Richard Adler	22b TELEPHONE (*Include Area Code*) 408-646-2352 22c. OFFICE SYMBOL 62Ab

DD FORM 1473, 84 MAR 83 APR edition may be used until exhausted
All other editions are obsolete

19. Abstract (cont'd)

3) HF communications are characterized by low capacity and high noise. A Logical Link layer to manage these is presented.

These three major thpics are integrated into a complete system using the ISO reference model. The following is achieved:

1) A fully integrated system across all frequencies.

2) ARPANET/DDN interconnect capability.

3) Efficient, effective HF links.

Approved for public release; distribution unlimited.

Ship-Shore Packet Switched Communications System

by

Rex A Buddenberg
Lieutenant Commander, United States Coast Guard
B.S., United States Coast Guard Academy, 1972

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN TELECOMMUNICATIONS SYSTEM MANAGEMENT

from the

NAVAL POSTGRADUATE SCHOOL
June 1986

ABSTRACT

This thesis presents an architecture for ship-shore sea service communications. Starting with the ARPANET packet switching model (TCP/IP), Network and Logical Link layers are defined which deal with the following problems:

1) Many ship-shore links are one way only. A Network level acknowledgement protocol to work under TCP/IP, integrates multiple, heterogeneous, one-way links into a complete network.

2) Conventional network access procedures are invalid in HF ship-shore communications because the assumption that ships can hear each other transmit cannot be made. Two network access techniques are presented.

3) HF communications are characterized by low capacity and high noise. A Logical Link layer to manage these is presented.

These three major topics are integrated into a complete system using the ISO reference model. The following is achieved:

1) A fully integrated system across all frequencies.

2) ARPANET/DDN interconnect capability.

3) Efficient, effective HF links.

TABLE OF CONTENTS

LIST OF FIGURES

Ship - Shore Packet Switched Communications System

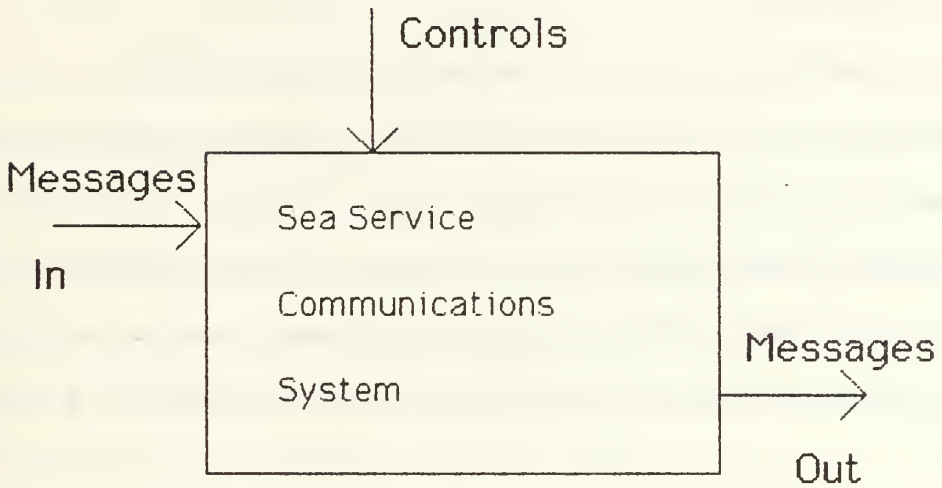

I. SHIP - SHORE PACKET SWITCHED COMMUNICATIONS SYSTEM

A. THE ENVIRONMENT

Computer to computer communications have advanced rapidly in the past decade. The trend is increasing as microcomputers proliferate and packet switching moves from the laboratory to industry. A great deal of study and organization has gone into the discipline culminating in efficient shoreside communications and a useful reference model (ISO) for decomposing the greater problem into parts that can be solved individually (See Appendix A).

Military communications, especially ship-shore communications, have lagged behind the technological and academic developments in packet switching. Part of this lag is due to the large existing investment that is difficult to change. Part of the problem is that the standards that have evolved--the X.25 standard, for example--are deficient when applied to sea service needs.

Specifically, High Frequency (HF) radio, as a communications technology, has not kept pace with the advances and has generally not been considered suitable for high speed data communications. Because of the difficulties of HF radio and the ease of use of the satellite media, the US Navy has allowed HF to languish. The equipment and procedures used today are the same as those of a decade ago and the architecture is thirty years old.

However, the communications satellite is not without problems. Those communication links may be vulnerable in wartime. Further, there is a limit to the bandwidth available,

while there seems to be no limit to the demand for communications, especially as off-board, over-the-horizon targeting becomes a reality. This demand for communications is burgeoning in the US Coast Guard as well, as new missions are increasingly more data intensive than the old ones.

B. THE PROBLEMS

This thesis examines and applies the advances in digital technology to the HF radio problem. Of primary concern are the issues of link integration, network access, and error control at the network and logical link layers in the HF system.

The problem examined is deliberately one of the more difficult. Long haul HF uses ionospheric sky wave refraction to get the necessary range. But the ionosphere contributes its own difficulties. The immediate primary problems are four. Two qualitative problems that make sea service communications different than conventional systems are:

1. Many physical links in sea service communications are one way only. Conventional networks are built on the assumption of full duplex communications--an assumption that is simply not the case for ship-shore links.
   A second aspect is that one ship may require greater capacity than is available on a single channel. The ability to harness several channels together is necessary.

2. All users cannot hear each other. This complicates the network access problem, especially in light of the constricted bandwidth.

Two are related quantitative problems:

3. A constricted bandwidth. A single HF narrowband channel will carry, at best, between 1000 (1K) and 10,000 (10K) baud. This is about 3 orders of magnitude less than the capacity of current conventional LANs.

4. A high noise level. One bit error in 10^2 or 10^3 is not uncommon. This is an error rate roughly three orders of magnitude worse than that experienced in conventional local area networks (LANs).

Additionally, there is a collection of characteristics imposed by the medium. This includes synchronous communications requirements (imposed by cryptographic and medium needs), and the frequency of communications entirely lost--fading (imposed by the medium).

Naturally, this environment can be worsened even further in a combat situation due to enemy jamming.

C. APPROACH

These four problems are dealt with in three major parts:

1. Lack of Full Duplex Capability

Many sea service communications systems are not duplex, nor can duplicity be faked by reversing the channel quickly. Many links are organized as fleet broadcasts and are one way only. For Emission Control (EMCON) reasons, as well as physical ones, this reality must be accepted.

This invalidates the conceptual communications model that underlays the CCITT and IEEE standards because they are predicated on the ability to provide full duplex communications at the physical level.

The second aspect to this constraint is the demand requirement to provide a user with more than one link--downward multiplexing. This is a halfway point between a local area net and an internetworking situation not seen in conventional systems.

The solution to the defective physical layer model is to recast the model with packet acknowledgement at the network level. This cannot be done with any existing standard. But once the concept of network level acknowledgement is accepted, solutions to the other related problems become readily apparent. Indeed the way now becomes clear to harnessing each communications band from ELF to EHF to work together in a fully integrated system rather than as separate systems.

A proposed Network protocol is presented to implement this solution.

2. Network Access

The second part of the problem is determining who talks when on a network. Only one station can talk at a time.

The ability of all subscribers to hear each other is an assumtion in collision avoidance, collision detection and token systems. In a long-haul HF environment, all subscribers must hear the network control station, but no assumptions can be made about subscribers hearing each other.

The only feasible solutions are centralized reservation systems. Since the Navy already has such systems in place, this solution is borrowed.

3. Efficiency in the Presence of Noise

Efficiency of use of the bandwidth available is necessary. Higher performance modems can increase the baud rate in an HF channel from the current 75 to somewhere in the 1-10k baud range. Data compression techniques can improve the throughput further.

But an increased data rate is only half of the quantitative picture. More error control tools are required to combat the noise (error rate) problem. The standard automatic repeat request technique of packet systems is necessary, but by itself, not sufficient.

Error control techniques must be implemented with care because they require bandwidth which is already a scarce commodity. In this thesis, three primary tools are implemented in adaptive fashion so that they impact bandwidth capacity only as necessary. Adaptiveness is required because the tradeoff between throughput and error control cannot be made at design time--the parameters change constantly.

This part of the thesis is the architectural design of link termination equipment specific to the HF environment.

D. RESULTS

The result of the thesis is an architecture that provides three major features:

1. Improved HF performance.
2. Packet switched interconnection capabilities--compatability with the ARPANET model.
3. Full integration across the radio spectrum.

E. THE HUMAN ASPECT

Finally, we need to understand that implementation of this communications system will change the role of the operator. Instead of sending and receiving messages, he will be observing process--the machinery will be performing the bulk of what is currently the radioman's workload. The radioman will become less a worker and more a supervisor.

The concluding remarks to the thesis cover the topics of frequency management and network organization. The attempt is to give the radioman and his communications officer an idea of what they can look forward to.

Link Integration -- Downward Multiplexing

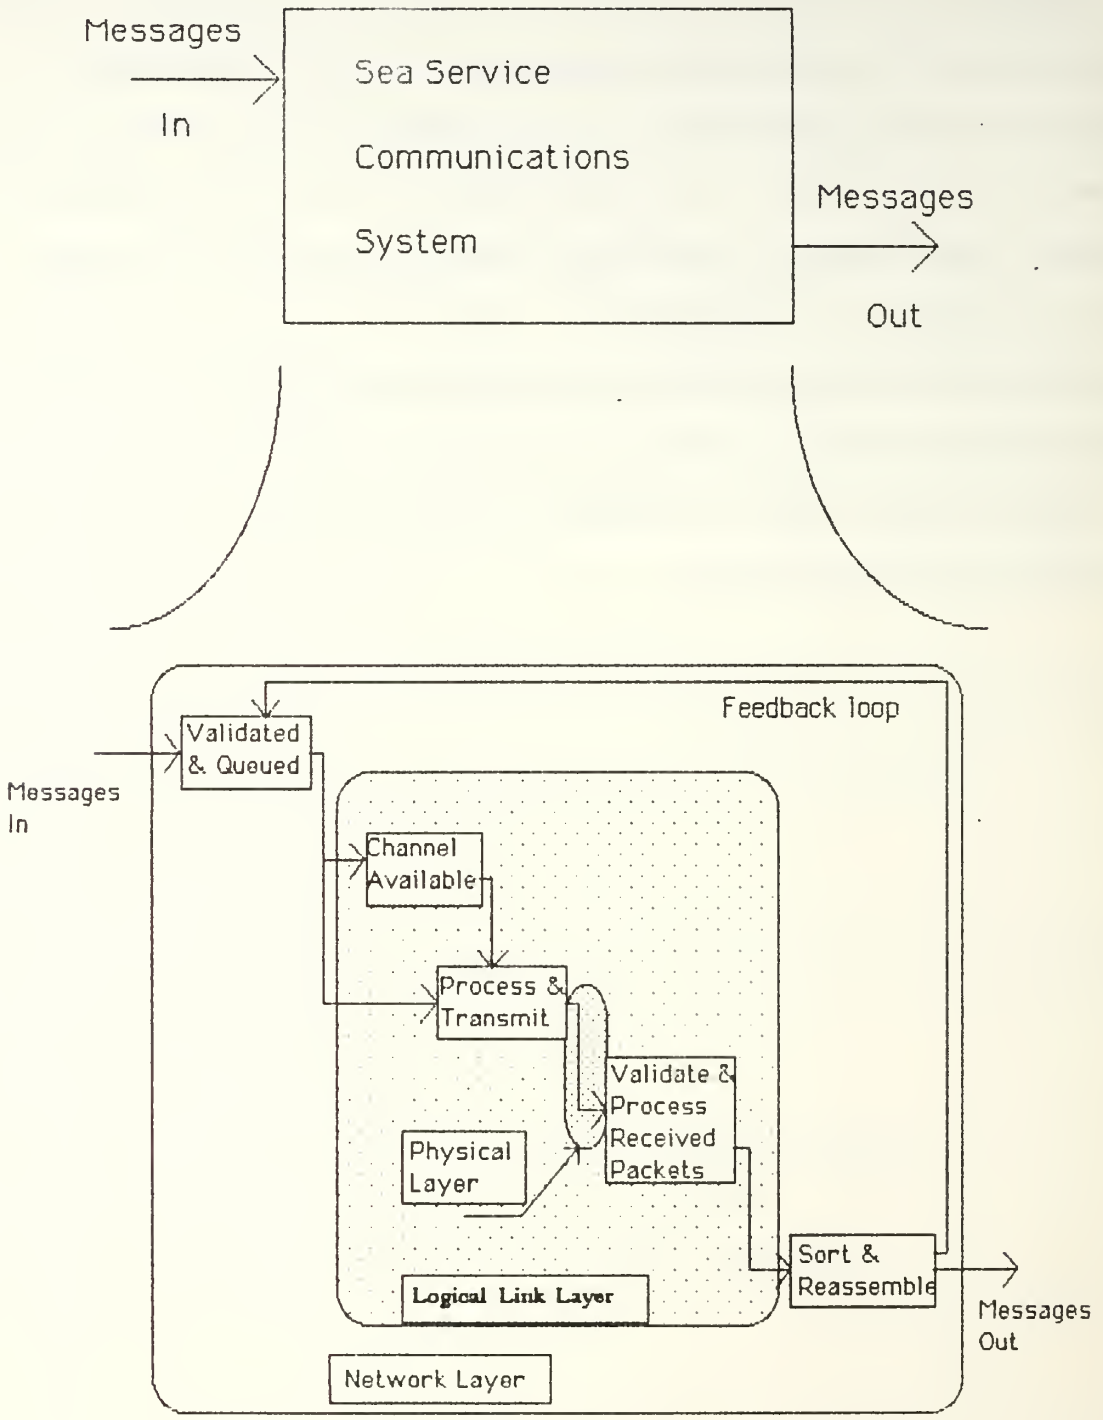

II. LINK INTEGRATION--DOWNWARD MULTIPLEXING

A. CHARACTERIZING THE PROBLEM

In this chapter, we view our ship-shore communications problem in a larger sense than just HF. The reason is that at bottom, there is a very important integration principle which makes ship-shore communications different in nature than shoreside communications. Failure to understand this difference would result in an inflexible and inefficient system.

The primary problem is that there is not enough capacity on any single link to satisfy the needs of a ship. No matter how much the capacity of a 3kHz channel is enhanced, there will be times when one channel is simply not enough.

The following facts about ship-shore communications in general and HF communications in particular affect the problem.

1. Factors that Affect the Issue

a. Efficiency onsiderations

(1) Ships can't full duplex. Because shipboard antennae are usually located close together, transmissions from a ship tend to drown out signals that the ship is attempting to receive.

This is complicated by the 'rusty bolt' effect where corners and dissimilar metal joints aboard ship tend to receive and rebroadcast transmissions on random frequencies. This problem is most pronounced in the HF band and makes it almost impossible to solve the drowning-out problem by filtering transmitters and receivers more carefully.

The wavelength of an HF signal ranges from about 10 to 150 meters; significant coupling between the antenna and the ship's structure takes place. Indeed, the ship's structure acts as part of the antenna. With this kind of activity taking place due to a powerful transmitter, it is little surprise that weak signals are often not received.

Finally, many radios are packaged as transceivers where the transmitter and receiver are in the same physical box. A switch couples one or the other to the antenna, never both.

The only real solution is a time division multiplexing one where the ship avoids transmitting when it is trying to receive.

Communications stations ashore generally do not have this problem because their transmitters and receivers are located several miles apart and communications stations lack the metal hardware that contributes to the rusty bolt effect.

(2) Attempts to emulate full duplex by line reversal are inefficient. Several items at the physical level require some preparation time before they are ready to transmit. High powered transmitters require a key-up time. Additionally, because of the often poor signal to noise ratios, sea service communications use synchronous communications, which requires a synchronizing preamble for the modem. Most cryptographic devices require a synchronizing preamble as well (the KG-84 is assumed for this thesis).

All this overhead is expended every time the ship shifts from receiving to transmitting and back. Packet systems tend to be 'bursty', which means that they will give up a great deal of efficiency to this overhead unless the bursty aspects are

confined to the data link and higher levels. At the physical level, the system must appear as continuous a stream as practical.

This subject is dealt with at length in the fourth chapter, but it bears on our problem here as well.

b. Capacity--Use of Multiple Channels.

This thesis assumes use of fixed bandwidth channels so the solution to a traffic load greater than the capacity of one channel is to allocate more channels. If a message is broken up into packets and a portion of these packets is sent over each allocated channel, the receiver must reassemble the packets into the complete message.

This solution is known as downward multiplexing and requires network and transport layers. How these layers are organized is critical to the network and will be considered in this chapter.

One characteristic of sea service communications is that of unbalanced traffic. It is not very common that a ship has the same volume of messages to go ashore at the same time that the communications station has messages to go to sea.

For tactical reasons, a ship may want to receive traffic, but will be unwilling to transmit (and send acknowledgements) at that time.

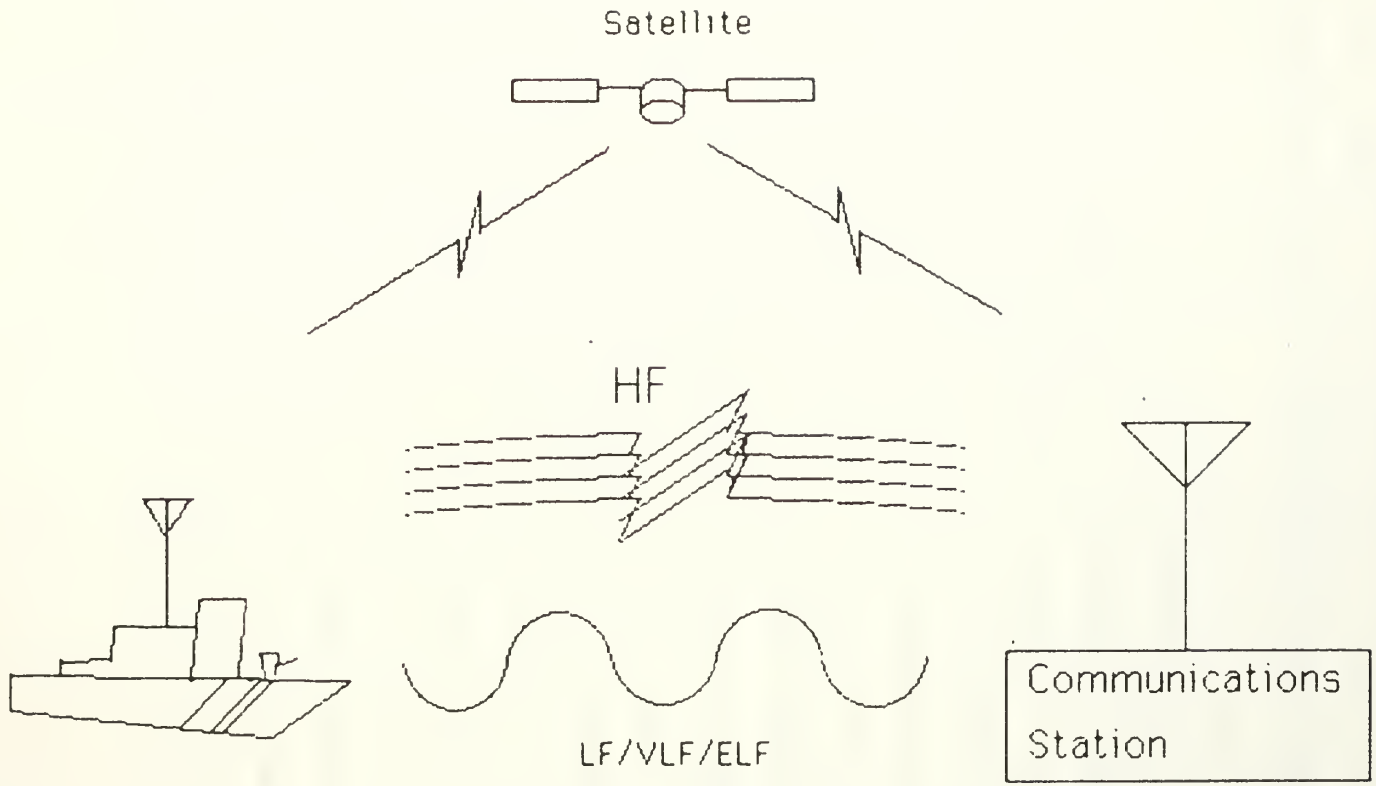

Downward Multiplexing

c. Electronic Support Measure (ESM) onsiderations

Because of its long range characteristics, HF signals carry their own special vulnerability. Radio signals in all bands have this problem, but they are the most severe in HF due to the large footprint caused by the worldwide propagation characteristics of HF.

(1) Direction finding vulnerability. An enemy who can intercept HF signals can, by use of cross bearings, locate a transmitter. Therefore ships may need to avoid transmiting in the HF band.

(2) Traffic flow analysis vulnerability. This involves an enemy intercepting signals and drawing conclusions about orders of battle and intentions by analyzing who is communicating with whom. Some value can be derived even if transmissions cannot be linked to specific transmitters. For this reason a decoupling of channels will be of value in frustrating this intelligence tactic.

For these reasons, we need the flexibility in our system that allows a ship to receive in one band (such as HF) and transmit in another (such as the satellite band) in an integrated fashion.

d. Fleet Broadcasts

The sea services organize some communications into fleet broadcasts. This is done for two primary reasons:

(1) Bulk. Below HF, transmitters require very large antennae and large amounts of power--impractical requirements for ships. A fleet broadcast in the VLF band, for instance, is one way only because of the physics involved.

(2) Current practice. In HF, two way communications is, of course, possible, but for efficiency and ESM reasons, a broadcast is frequently used. It should be noted that the Navy operates a broadcast in virtually every band that it uses for ship-shore communications. Because of these efficiencies, we need to integrate the concept into our network.

We can divide communications into three classes, based on how the messages are accounted for:

1. Full_ARQ. This type of message is receipted for by the receiver. Within this thesis, we will refer to these as full_ARQ (full Automatic Repeat reQuest). Delivery of a full_ARQ message generates an acknowledgement known as a QSL in radioman terminology. Incorrect receipt requires a retransmission by the sender which is initiated by a NAK, a Negative AcKnowledgement, also known as a ZDK. Primary ship-shore systems use this method which assumes delivery only when the sender gets an acknowledgement (ACK). (See glossary for operator signals and acronym decoding.)

2. NAK_only. A Negative AcKnowledgement only method is characterized by the current fleet broadcasts where messages are assumed delivered when sent. Retransmissions occur only upon receipt of a NAK by the sender. Messages are serially numbered so a receiver can detect missing messages.

3. No_ACK. Non-acknowledged messages are currently more common in intra-task force communications than in the ship-shore portion of the sea service system. Examples of content include track reports where it is less efficient to wait for a refresh of the data than it is to seek a retransmission.

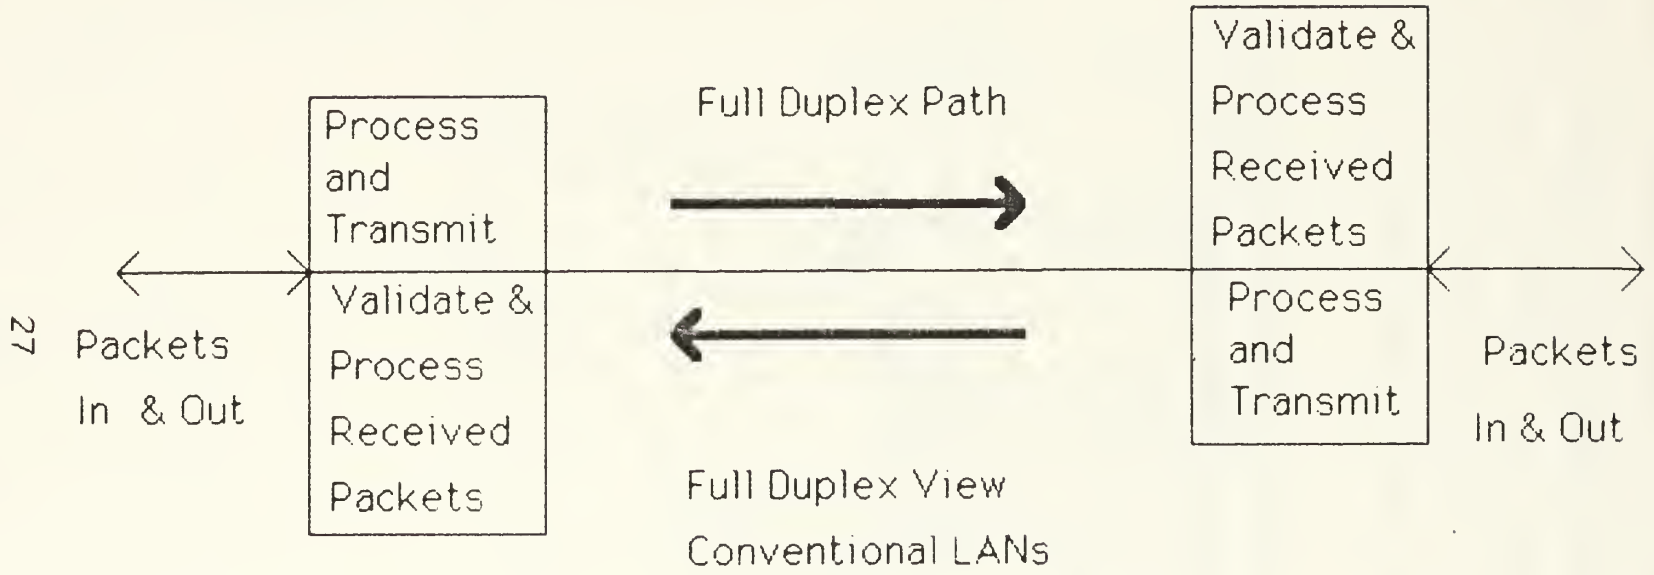

Conventional Physical Link Model

2. Conceptualizing Conventional Networks

The conceptual model of conventional communications is that of a full duplex channel at the physical level. That feature is unfortunately not isolated to the physical layer, as is the partitioning intent of the ISO model, but affects the structure of the link and network layers in the existing standards.

For instance, using our packet terminology, a Packet Assembler / Disassembler exists at each end of the link. The feedback loop (ACKnowledgements) is simply the reverse of the information channel.

More pervasive is the structure of the current standards. The X.25 standard, which is typical in this regard, uses a link level acknowledgement system that can be illustrated using our SADT technique.

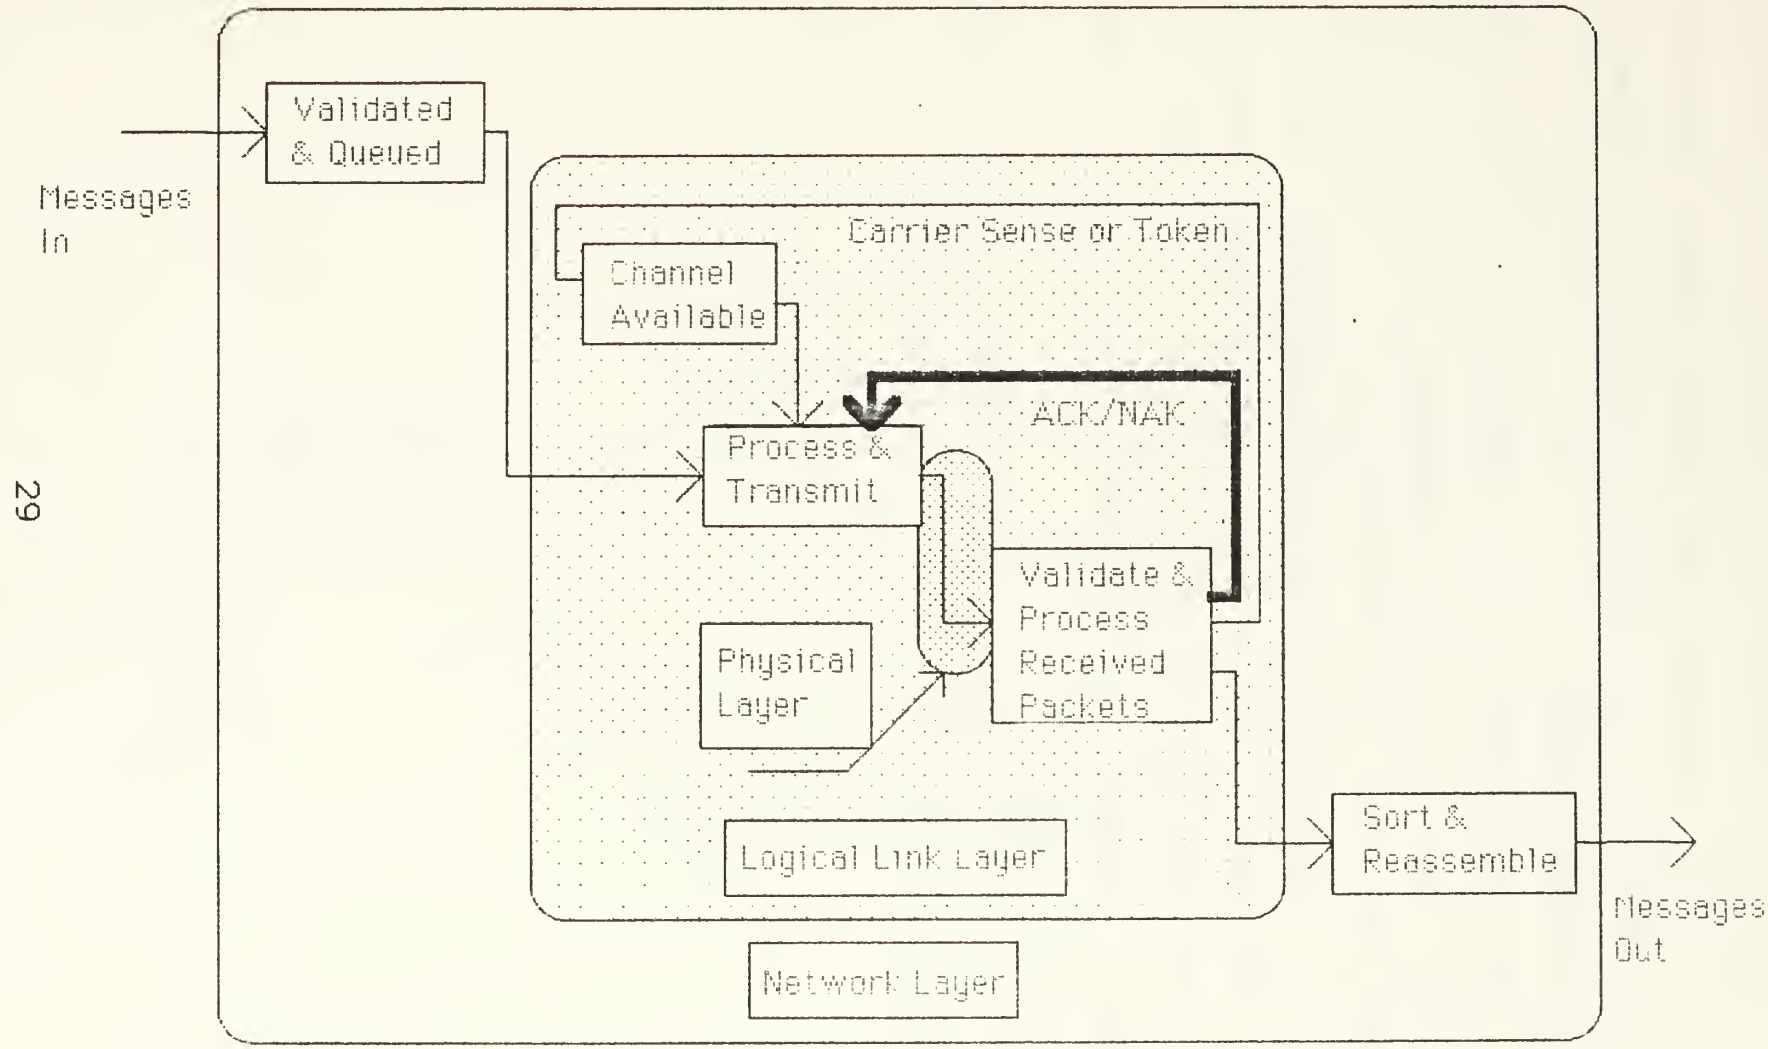

Conventional Local Area Network

The needs enumerated at the beginning of this chapter are not met by this approach. Our desired goals of downward multiplexing and use of fleet broadcasts in not compatible with this conceptual model.

The flaw resides at the fundamental conceptual model of a physical link. Sea service communications links must be modeled as one way links, as illustrated on the following page.

3. Conceptual Foundation for Ship-Shore Links

This one-way channel model is by no means incompatible with full duplex systems--if the physical layer is indeed full duplex, we simply model that as two of our one-way links--one in each direction.

By use of this model, we are now able to build a system that meets most of the needs that we have specified.

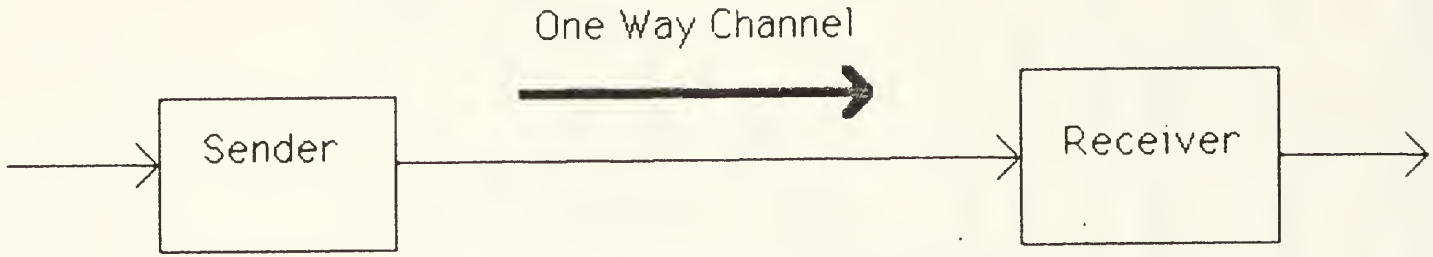

Seq Service Conceptual

Physical Link Model

At the level of protocols to handle packets, the basic model must change to handle this problem. Conceptually, this requires a promotion of the acknowledgement system from the logical link to the network level. In this way, the feedback loop remains intact, but it has been decoupled from the specific physical channel--it is only required that the receiver have any existing channel to the sender, not a specific return over the same channel that it has received data.

This can be represented by the following diagram. Acknowledgement packets are managed just like information packets, as far as the link level is concerned.

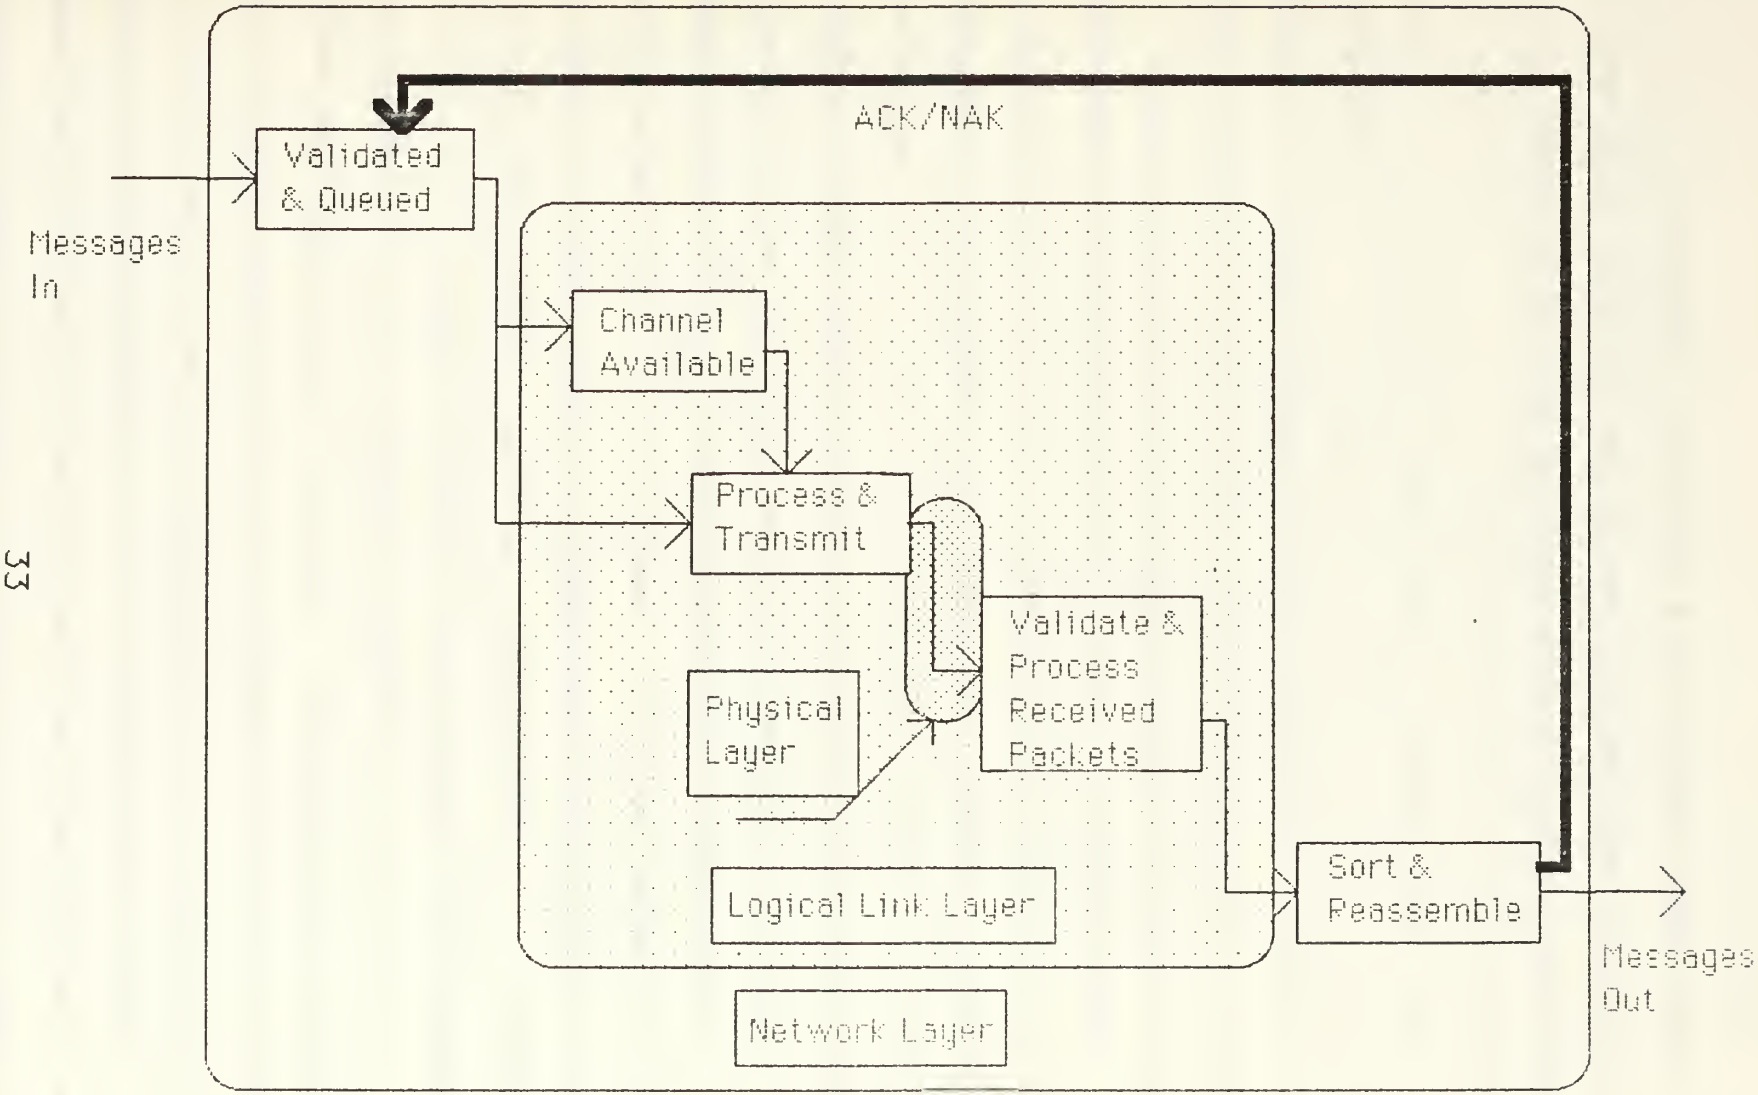

Sea Service Networ

B. DISADVANTAGES

Since the reader should have some sense of the advantages to be gained (because the problems to be addressed were enumerated at the beginning of the chapter), we will examine the implications and difficulties first.

The first problem is that there is no protocol in existence that handles network level acknowledgement. Virtually every conventional network provides full duplex service at the physical level, so there is no perceived need until the problems of sea service communications are mapped onto packet switching.

This problem of protocol inadequacy must be addressed early and directly. If it is postponed, it will become increasingly expensive to correct. This has been demonstrated repeatedly in software engineering projects where design errors became evident at the coding and testing stages of projects [Appendix D].

The second problem with this conceptual model is that it will require more computational capability than those standards that are supported by full duplex links. This problem is not severe as the requisite capability exists, and it is only large compared to existing systems, not when compared to the microprocessor capabilities available today.

In particular, a robust transport layer is required to detect missing messages. The network acknowledged ship-shore system is a classic example of the 'unreliable' network where missing packets are not detectable by the network itself.

Packet fragmentation is a third problem that derives from the bandwidth and noise available within the physical channels. It is now entirely feasible for a complete message to be sent on

several different physical channels with varying packet sizes. There is a large amount of research in the ARPANET literature addressing essentially this problem, so the problem is neither new nor insoluble.

C. ADVANTAGES

The first obvious advantage is the ability to easily downward multiplex. It is now practical to transport a large amount of data from a sender to a receiver by using several one way channels. If the receiver only has the acknowledgements to send, it may only require one channel to return them. So the system is fully adaptable to the amount of traffic offered in each direction.

Since the send and receive physical channels are decoupled, the ESM problem becomes somewhat more tractable. A ship can use a channel different from the receive channel--whatever is tactically wisest--to close the loop. This means that HF, VLF, and satellite channels can be used complementarily rather than separately. This potential for synergism appears to be a great possibility for improving the use of the spectrum--getting more communications out of the existing bandwidth.

Fleet and tactical broadcasts are fully included in the model. By identifying an individual packet as full_ARQ, NAK_only or no_ACK all three classes of messages can be managed--if necessary, on a single channel. This requires some work at the protocol design stage to include a mechanism for packet accountability (similar to the existing message serial numbering

system) so that the receiver will know when it missed a packet. This is part of the Network Protocol issue of the next chapter. Once the design work is done, the same hardware and software that supports a fully interactive network (primary ship-shore) can be reorganized to run a fleet broadcast type of system. This becomes a function of the message content, not the system design.

D. IMPLICATIONS

There are two fundamental implications of this network level acknowledgement structure. The first is a decoupling of the send and receive CHANNELS as discussed above.

The second is the TEMPORAL decoupling of an information packet and its corresponding acknowledgement packet. There is no longer any imperative of the communications system itself that requires a packet to be acknowledged immediately. A packet received can be acknowledged in a time frame controlled by the precedence and content of the packet itself, not by the communications system.

This temporal decoupling is important for ships operating in a tactical environment where it may not be prudent (or possible -- as in the case of submerged submarines) to respond immediately.

E. CONCLUSION

Sea service communications IS different from conventional shoreside communications. The fundamental difference is that ship-shore communications must be modeled as a network of one-way physical links, not as full duplex ones. This requires a network level acknowledgement, rather than a logical link level one.

The next chapter outlines a Network Protocol necessary to implement a network level acknowledgement system.

Network Protocol

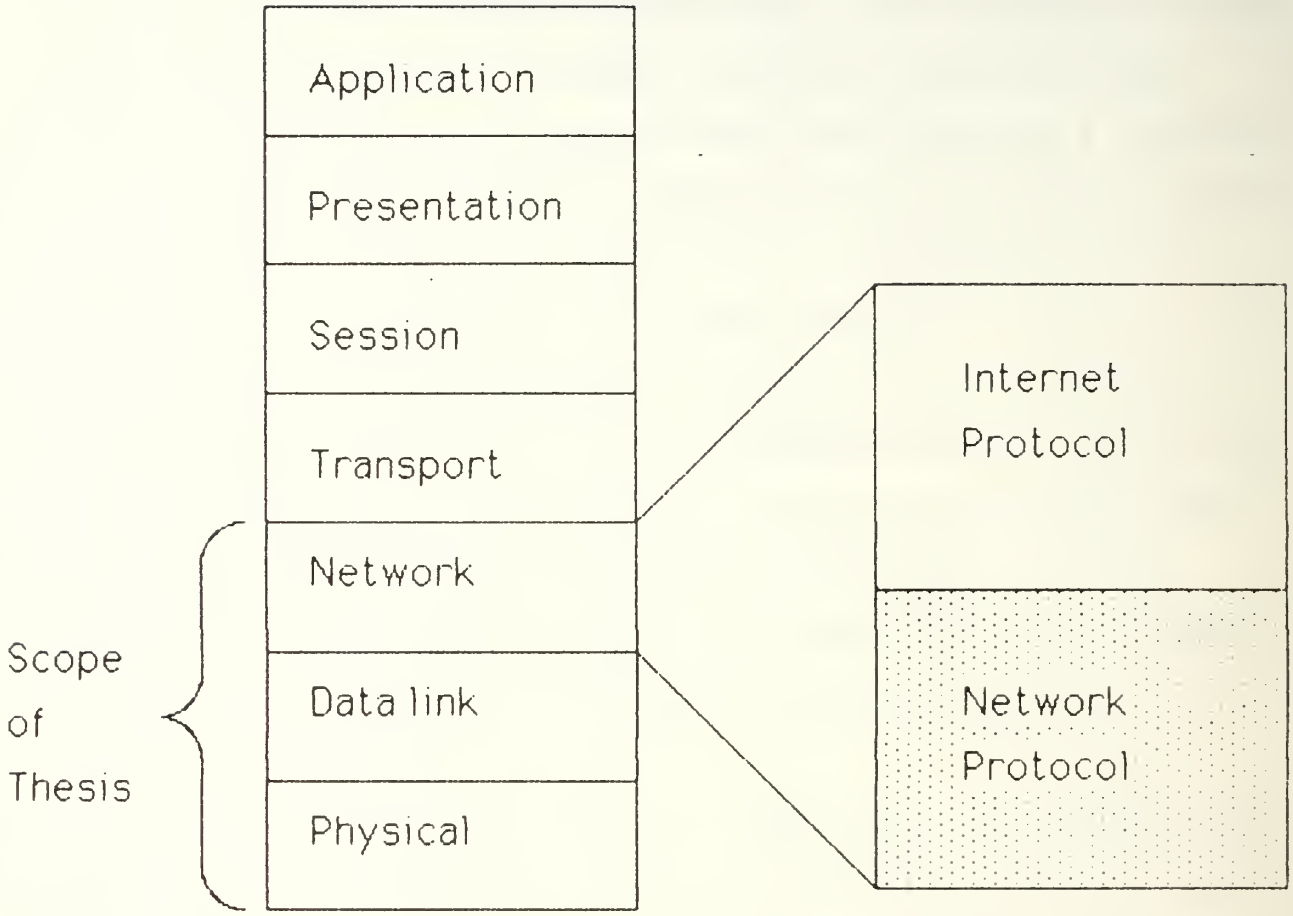

III. A PROTOCOL FOR NETWORK LEVEL ACKNOWLEDGEMENT

A. OBJECTIVE

The purpose of the previous chapter was to outline the need for a network layer acknowledgement system and to lay the conceptual foundations. The objective of this chapter is to set forth a draft protocol to accomplish this network layer acknowledgement. The Network Protocol must weld multiple, heterogeneous communications links into complete networks. Therefore it must be general enough to handle any specific link, be it HF, satellite, VLF or LF, or shoreside telephone.

A well understood interface between the network and logical link layers is mandatory to successful decomposition of the software problem and the key to maintaning an open, expandable architecture.

1. Two Points of Departure are Available

a. Internet Protocol

Internet Protocol provides a network layer (upper half) protocol for operating across multiple networks. As such it is not directly concerned with intra-network communications. Most functions that are handled within Internet Protocol will not be duplicated by the intra-network protocol. The two must mesh cleanly.

Internet Protocol can be visualized as the upper half of the network layer. For clarity, we will use the term Network Protocol for our lower half, intra-net construction to avoid confusion with Internet Protocol (IP).

b. Local Area Network Protocols

X.25 is one of the best known intra-network protocols and is representative of those available. It is also the protocol of choice for connection to Arpanet. Therefore, we will adapt what can be drawn from X.25 into the Network Protocol.

Unfortunately, X.25 and related protocols all depend heavily on the full duplex physical layer model. Consequently, they cannot simply be modified for our purposes.

2. References

Since we will not discuss the existing standards here, the reader may wish to consult the references for sources of each [CCITT, 84; DDN-X.25, 83].

B. PACKET HEADER

Demands upon a Network Protocol structure--the packet header. Packets, when delivered from logical link machines, should contain the following information. Some of these parameters must physically be transmitted in the header of each packet. Others may be supplied by the receiver in its processing activity. For instance, the channel identification--the frequency used--is certainly known by the receiver and needn't be transmitted with each packet using up scarce bandwidth.

On the other hand, we needn't be overly concerned about size of the packet header. While bandwidth conservation is important, room for growth in the protocol is needed and the bits expended here can be regained through data compression at the link level.

1. Addressing

X.25 uses a logical address for each subscriber. A drawback is that a message addressed to multiple subscribers must be sent individually to each. Since a great deal of sea service message traffic is addressed to multiple destinations, an addressing scheme that allows multiple addressing will greatly improve efficiency. This is important in our bandwidth constricted environment. Use of the existing system of call signs, Address Indicating Groups and Collective Address Designators is quite usable.

The objection that such a scheme uses a significant number of bits can be overcome by use of a data compression algorithm at the logical link level. Since the bit patterns of addresses are quite predictable, efficient compression is practical.

A collective addressing system allows a communications station to send a packet but once, then collect all the acknowledgements. Obviously, retransmissions may be required to overcome link unreliability and noise, but the number of retransmissions will be much less than if each ship were serviced individually.

2. Message Identification

Several points of message identification are required. This section is necessary for two reasons. The first is that of fragmentation. Packets must be sized to the needs of the communications channel. These needs, and consequently packet sizes, vary from channel to channel and moment to moment. The Network and Internet Protocols must be capable of reassembling messages, parts of which are sent over widely varying paths.

The second reason for careful identification is to control duplicates. Shore to ship circuits especially, are often redundant in order to increase the probability that a message will be received at least once. One example includes keying multiple frequencies (QLH) in the HF fleet broadcasts. A second is queuing up messages destined for submarines on both satellite and VLF channels.

A packet switched system should be able to handle two related cases:

First, if a message is received over multiple paths, the receiver will have several duplicate copies. These must be recognized and the duplicates eliminated. Further, a single acknowledgement should be recognized by the sender as adequate to control all the queues, not just a single channel.

Second, if parts of the message are received on one channel and parts on another, the system must be able to merge them correctly into a complete whole. This need is occasioned by two factors. One is the downward multiplexing circumstance where parts of the message are sent over multiple channels in order to increase throughput. The second is channel unreliability--a receiver may get fragments of a message on several channels. The ability to synthesize these fragments into a complete whole is valuble.

a. Channel Identification

Identifying the channel over which a packet was sent is necessary because the condition of the packet on arrival is an indication of the quality of the channel. If some adjustment is required, the specific channel used must be known.

b. Message Identification

This is analogous to the Date-Time Group of current messages. There is some capability for identification in the Internet Protocol, so it may only be necessary to copy the data from the Internet Protocol header to the Network Protocol header.

c. Packet Identification.

This is the information necessary for gluing packets back together into complete messages. This is also part of the Internet Protocol and should not need to be duplicated.

d. Packet Offset

When packets are broken down into smaller fragments to meet the requirements of the channel, some control information to reverse this fragmentation will be necessary. This will become particularly important as the network grows to encompass several diverse types of links, all using different packet sizes. The packet identification capability of Internet Protocol is not adequate by itself.

When assembling an irregular collection of packets into a complete message, the assembler can map efficiently if it knows how far from the beginning of the message this packet starts. This overlaying can be effective when packets derived from one channel overlap those from another.

Packet offset and message identification together can be used to patch together a message received over several noisy channels. An acknowledgement of the whole message (ACK by msg_id vice packet_id) should cause the sender to cease retransmissions of any specific faulty packets (parts) of the message.

Packets from different noisy channels
Overlaid to form a correct, complete message.

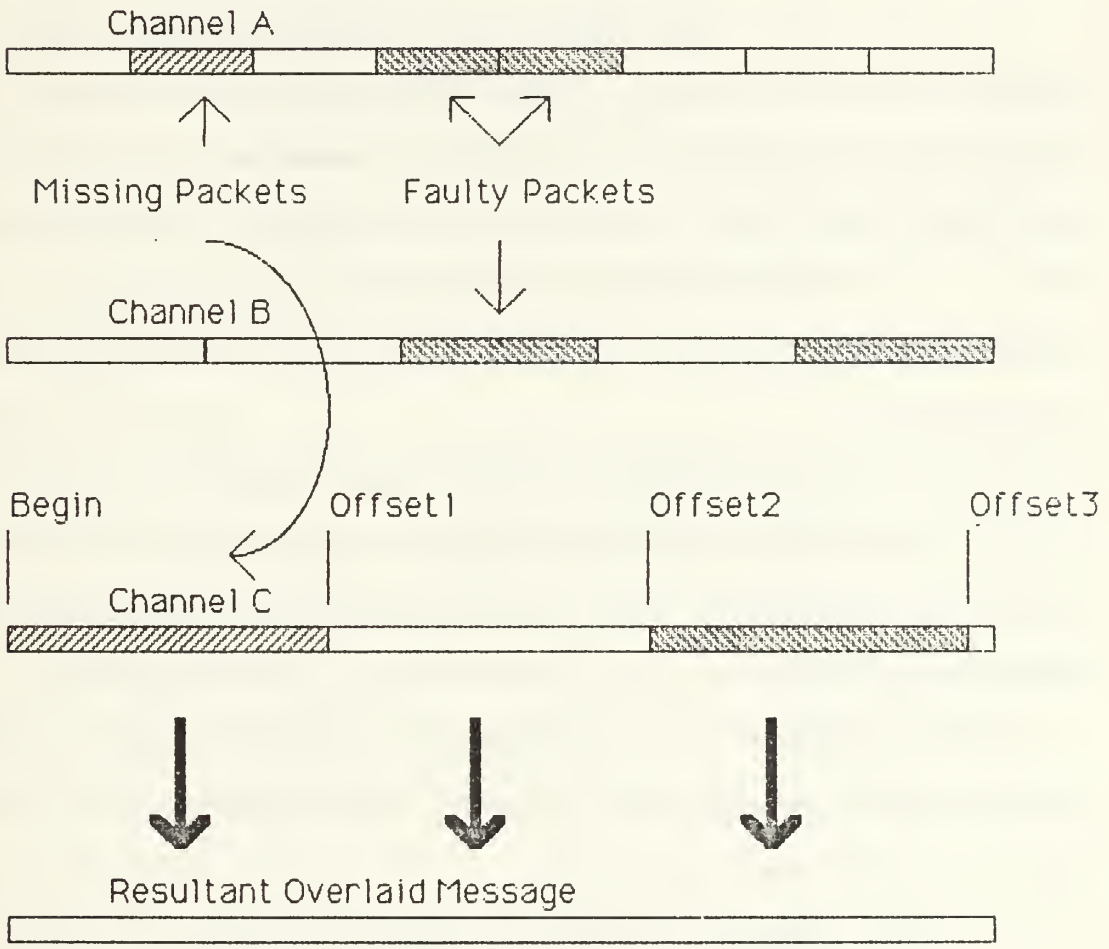

3. Control Information

This is information telling the receiver how to handle particular packets.

a. Acknowledgement equirements

(1) Full ARQ. This is the only type of packet represented in conventional networks. Every packet generates either an ACK or NAK from the receiver.

(2) NAK only. This is the type of packet epitomized by the current fleet broadcast where receivers are passive and only service for missing packets. All packets are considered delivered, unless otherwise noted. This requires some care in the construction of message identification because the transport layer must be able to recognize a missing message circumstance.

(3) No ACK. This type of packet is illustrated in the NTDS Link 11 system. Because this kind of data is frequently refreshed, it is more practical to simply wait for the updated data than to try and retrieve a flawed packet.

Since Internet Protocol does not address these issues and assumes that all packets are to be acknowledged, we must include them in our Network Protocol.

b. Degree of Error Freedom Required

Some data, such as computer programs and supply data such as MILSTRIPS must be fully error free. Others such as graphics or voice can tolerate more errors without losing their usability.

This is treated within Internet Protocol, but probably not adequately. Internet Protocol gives only two alternatives:

normal and high reliability. It is beyond the scope of this thesis to haggle over whether Internet Protocol can or should be modified; suffice it to say, that we need the differentiation.

c. Precedence--Grade of Service Indicator

This is adequately addressed in Internet Protocol and only needs to be copied into the Network Protocol header.

d. Time to Live

Some data, such as the NTDS No_ack type illustrated above, and also meteorological reports, are highly perishable--like yesterday's newspaper. If they are not delivered within a certain time frame, they should be removed from the communications system to ease congestion and allow more recent information through.

In the case of NTDS data, outdated packets should simply be thrown away. Meteorological data have climatological significance which indicates that delivery is useful, but not at the original precedence. In this case, the precedence should be downgraded, possibly to the point where the message is removed from the active queue and spooled onto tape where it can be mailed to the destination.

Internet Protocol contains a time to live parameter but it is used for a different purpose--to eliminate 'stray' datagrams (messages) that might otherwise exist on the net forever. Since our purpose is congestion control and maintaining a grade of service in the presense of high loads, the Internet Protocol approach must be modified.

Some packets should be given a time to live measured in seconds. These can be handled adequately within the Internet Protocol specification. But this function should be handled as low in the ISO layering as practical. The Network Protocol is the first practical layer to handle the problem.

e. Version Number

Including a version number in the packet header allows a receiver to tell if this packet was generated by a Model A or Model B sender. This allows for incremental updating of the link equipment as ships return to port. Incremental updating avoids the awkward and difficult situation where it is necessary to update all the units in a communications area simultaneously.

f. Header Checksum

A header checksum allows a receiver to check for errors in the header that might result in packets being routed to the wrong destinations.

C. INTERFACES

It is important to understand clearly how each layer communicates with superior and subordinate layers.

1. Network to Internetwork Interface

This joining is the simpler of the two. The Network Protocol passes packets to the Internet Protocol, just as existing network layers (e.g. X.25) do. The implication is that the Network Layer, upper half, and all higher layers will not require modification.

2. Network to Logical Link Layer Interface

This interface is more complex. First, let's analyze the types of packets that will be present at the network layer.

a. Data Packets

Data packets are delivered to the Internet Protocol, as is noted above.

b. Control Packets

Control packets, including all ACKs, are used by the Network Protocol and lower layers and are not delivered to Internet Protocol. These are the materials that make the network level acknowledgement function. These packets will not be visible to the Internet Protocol or higher layers.

(1) Acknowledgement packets carry the same precedence as the data packets that caused them. This results in a temporal matching between the precedence of acknowledgements based on the content of the corresponding messages, not on the exigencies of the communications system.

The receiver must discriminate between them and treat ACK packets as higher precedence than equally marked data packets. This is needed to preclude the circumstance where a data packet keeps timing out and being retransmitted and preempting the ACK necessary to turn off the retransmissions.

(2) Other control packets include those necessary for passing link parameters such as Sequence Order Lists or transmission cues, error rate, packet length requests, coding and baud rate and circuit logins. These are necessary for network control purposes, and normally will not be passed above the Network Protocol boundary. Most of these packets, because they

are vital to maintaining communications, will be the highest precedence packets in the system.

3. Multiple Ports

Since the function of the Network Protocol is to join several logical links representing several types of physical media together, there must be multiple ports to the Network Protocol. Each port will have two components:

a. Packet Port

The packet port will be the port where data packets are either sent or received. Service will be in a queue fashion where the highest precedence packets stand at the head of the queue.

b. Control Port

The control port will be a bidirectional port where the control packets are sent and received. Control packets cannot wait in the queue; they must be delivered to the sender or receiver as soon as the Network Protocol receives them.

The upshot of this Network Protocol is that several functions are performed by it that were formerly done by the logical link layer. This simplifies logical link design, which is beneficial because there are several efficiency problems that must be dealt with there, that are not as severe in conventional networks as in our ship-shore system.

Before departing from the Network Protocol and launching into a detailed discussion of the components of our link termination devices--logical link components--we must esablish a firm interface between the network layer (node) and logical link layer (link).

One major reason is an engineering one. It makes a great deal of sense to divide a development project at this point and engineer nodes and links separately. This is particularly important if we wish to engineer several different (e.g. HF, VLF and satellite) links.

Initially we must consider the data stream. Each link will present or demand data at different and varying rates. For instance, in our HF system, packet size, compression efficiency, error code rate and modem baud rate will all affect the input data rate to a sender. Additionally, the requirement to retransmit packets will affect the input data rate.

Similarly, a VLF broadcast will demand data at a different rate than an HF one.

Since our link and physical level equipment function best with continuous rather than bursty data, the data stream must be capable of supplying bits as fast as they can be transmitted.

Receivers will present data to the network at different and varying rates as well. For instance, if a receiver is momentarily unable to receive correct packets due to a high error rate, no data will emerge (save NAK packets) until the retransmissions catch up.

Therefore the node equipment must be adequately buffered to accept and offer data with the speed and immediacy required by

the link termination equipment. Since the ship-shore data rates for each link are modest by comparison with shoreside LANs this should not be a problem. Flow control may be required for the shoreside links, but the system should be sized so the ship-shore links will not need it.

The data stream itself should be managed as a queue. The node equipment must manage a queue for each sender. This queue is maintained by the network layer in precedence order. Network level acknowledging allows the ACK/NAK packets be treated the same as information packets. An Immediate precedence information packet should generate an Immediate precedence ACK when received (assuming full_ARQ). ACK/NAK packets of a particular precedence should be queued ahead of information packets of the same precedence.

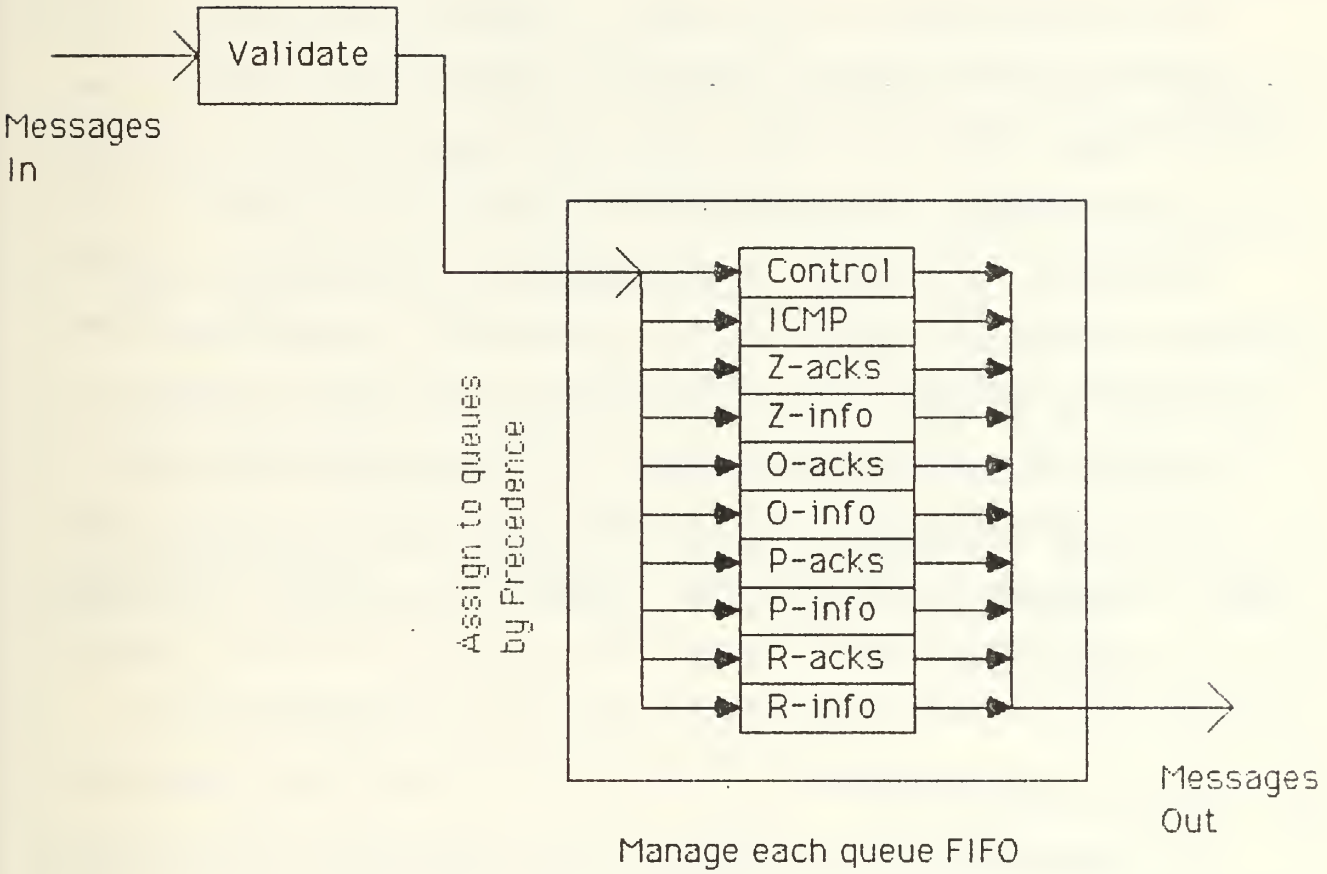

Queue Management

Since the node will not be aware of packet size within a link, it can only queue up a data stream (which may consist of several packets or messages). This data stream should consist of messages preceded by the header of the message. When the sender is ready to send a packet, it reads in the required number of bits from the queue and replaces the header at the head of the queue.

This allows the link and node to function independently--the node may need to queue higher priority traffic ahead of the existing queue contents. The queue is pushed down and the more imperative message is inserted at the head of the queue.

Meanwhile, the sender is processing and sending the bits that it has lopped off the front of the queue, unaware that a change may be simultaneously occurring to the queue contents.

A break must be placed between each message in the queue. If the sender is reading in bits to be sent and encounters the break (in ARPANET this is called a push), it chops off the packet at the break point.

D. CONTROL INFORMATION

The second port in the node-to-link interface is the control information port.

This control information that the packet machine needs cannot be passed through the same port as the data stream. Control packets must go to the link control unit by a separate channel. Examples of this control information include the contents of the Sequence Order List or cuing signals (so the sender knows when to transmit), and control packets originated by

a receiver requesting a change in packet size, error code rate or bit rate.

An additional control packet would be one from a station logging into or out of a link. This would require actions of the control unit such as alloting space in the SOL or returning sent but unacknowledged packets to the node for rerouting to another link.

In addition, if more than one HF transmitter are ganged together, the two link senders must be able to pass coordination messages back and forth. This kind of control includes a receiver inhibiting a sender while it is actively receiving traffic. This is information that is needed by adjacent links to synchronize themselves--non-adjacent links shouldn't need it. But this information is of no value to the node. Therefore, we need a set of multiplexing control units from adjacent links connected together. The control messages can be passed via the network, since the control ports must be bidirectional.

In the sections below, we discuss how these control packets and queues are managed in a complete system.

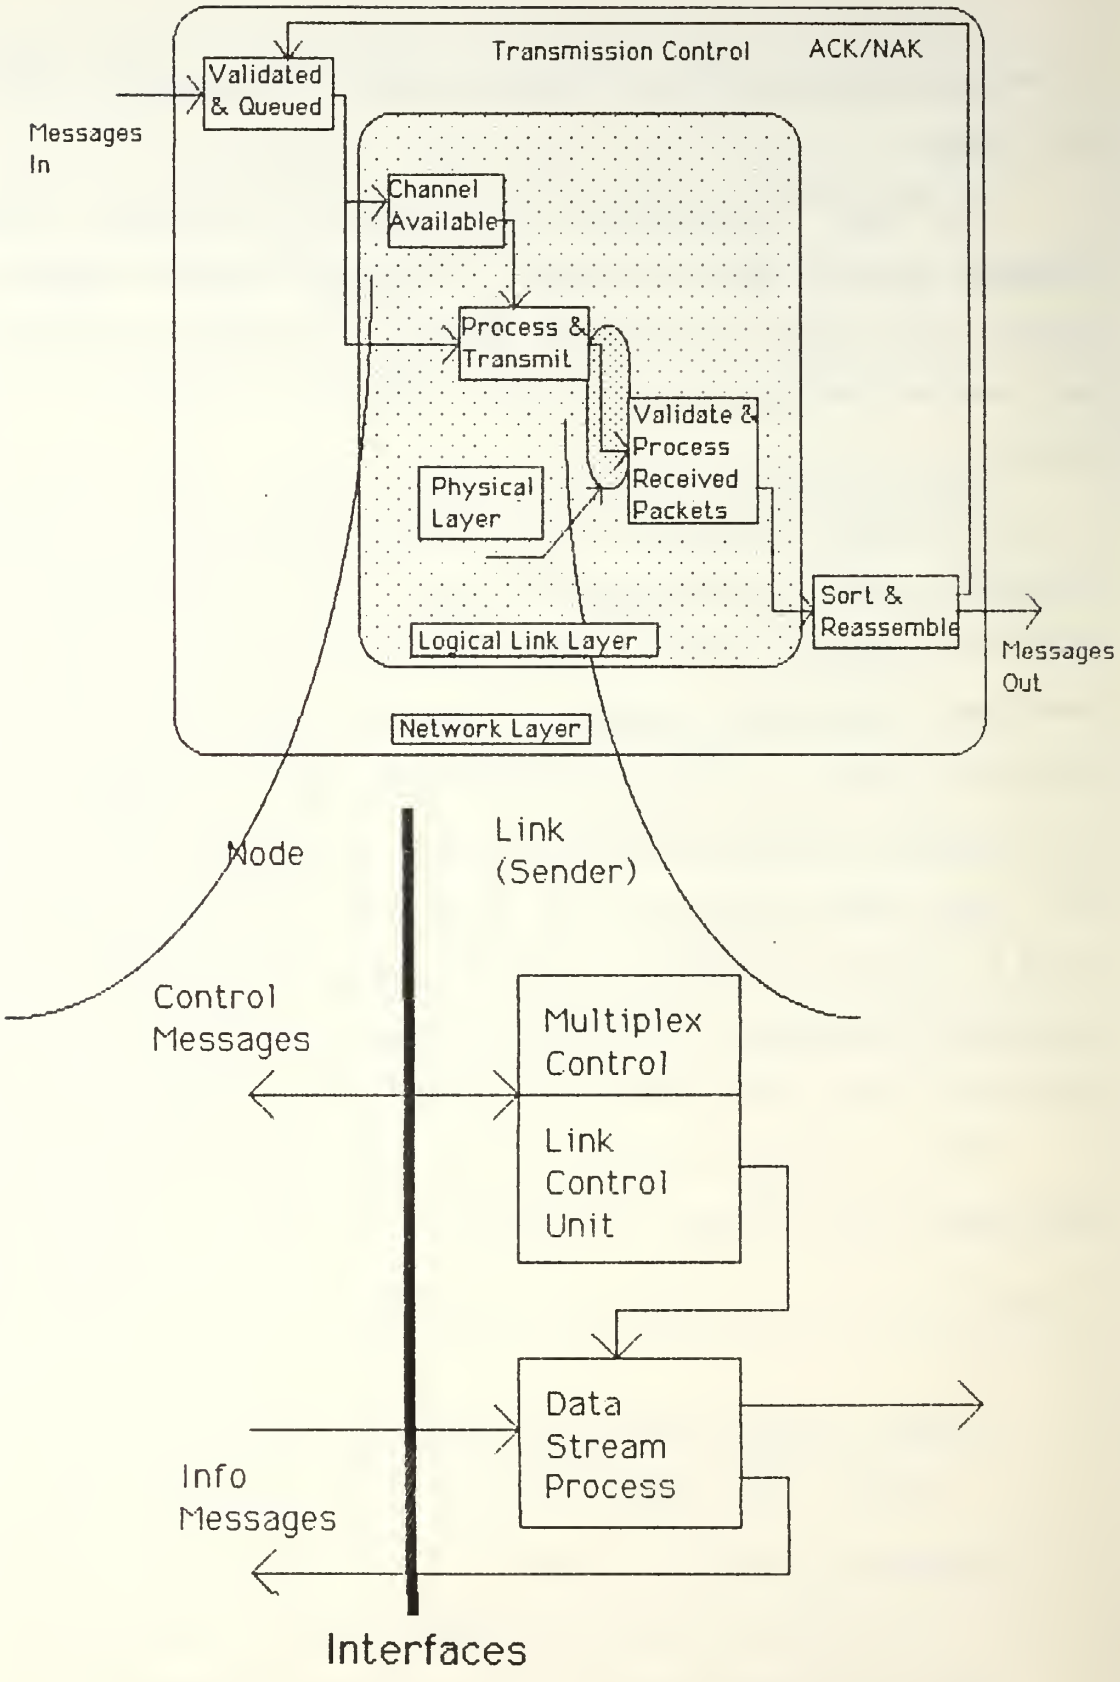
Transmission Control
ACK/NAK
Validated & Queued
Messages In
Channel Available
Process & Transmit
Validate & Process Received Packets
Physical Layer
Sort & Reassemble
Messages Out
Logical Link Layer
Network Layer
Node
Link (Sender)
Control Messages
Multiplex Control
Link Control Unit
Data Stream Process
Info Messages
Interfaces

In summary, a logical port between a node and a sender will need two bidirectional physical ports:

1. Node delivers data to sender (queue) and sender returns undeliverable packets to node
2. Node delivers control information to sender and one sender coordinates with another.

A logical receive port will consist of two physical connections:

1. A receiver to node data port which will carry output packets. These packets, in turn, fall into two categories:
   - 1a. Information packets to be passed through to the Internet Protocol
   - 1b. ACK/NAK packets to be routed back to the sender by the Network Protocol
2. A control port to provide:
   - 2a. A channel to deliver control packets to the node for routing to the opposite node. This provides a way to control physical parameters and to log into and out of networks.
   - 2b. A node-to-receiver control line, to force the output of pending (errored) packets.

E. STATE INFORMATION IN A NETWORK PROTOCOL

This information is not sent with the packet, but rather stored with the packet in a pending area (waiting queue).

1. Retransmission Timer

Known as T1 in X.25, this is the amount of time that elapses before an unacknowledged packet is resent.

When the packet is acknowledged, the packet is removed from the waiting queue and logged on archive as sent and acknowledged.

If T1 times out before an acknowledgement is received, the packet moves to the send queue and is resent. (Note that this

retransmission may or may not be over the same channel that the original packet was sent.) It is then returned to the waiting queue with a new T1.

Packet State Transitions

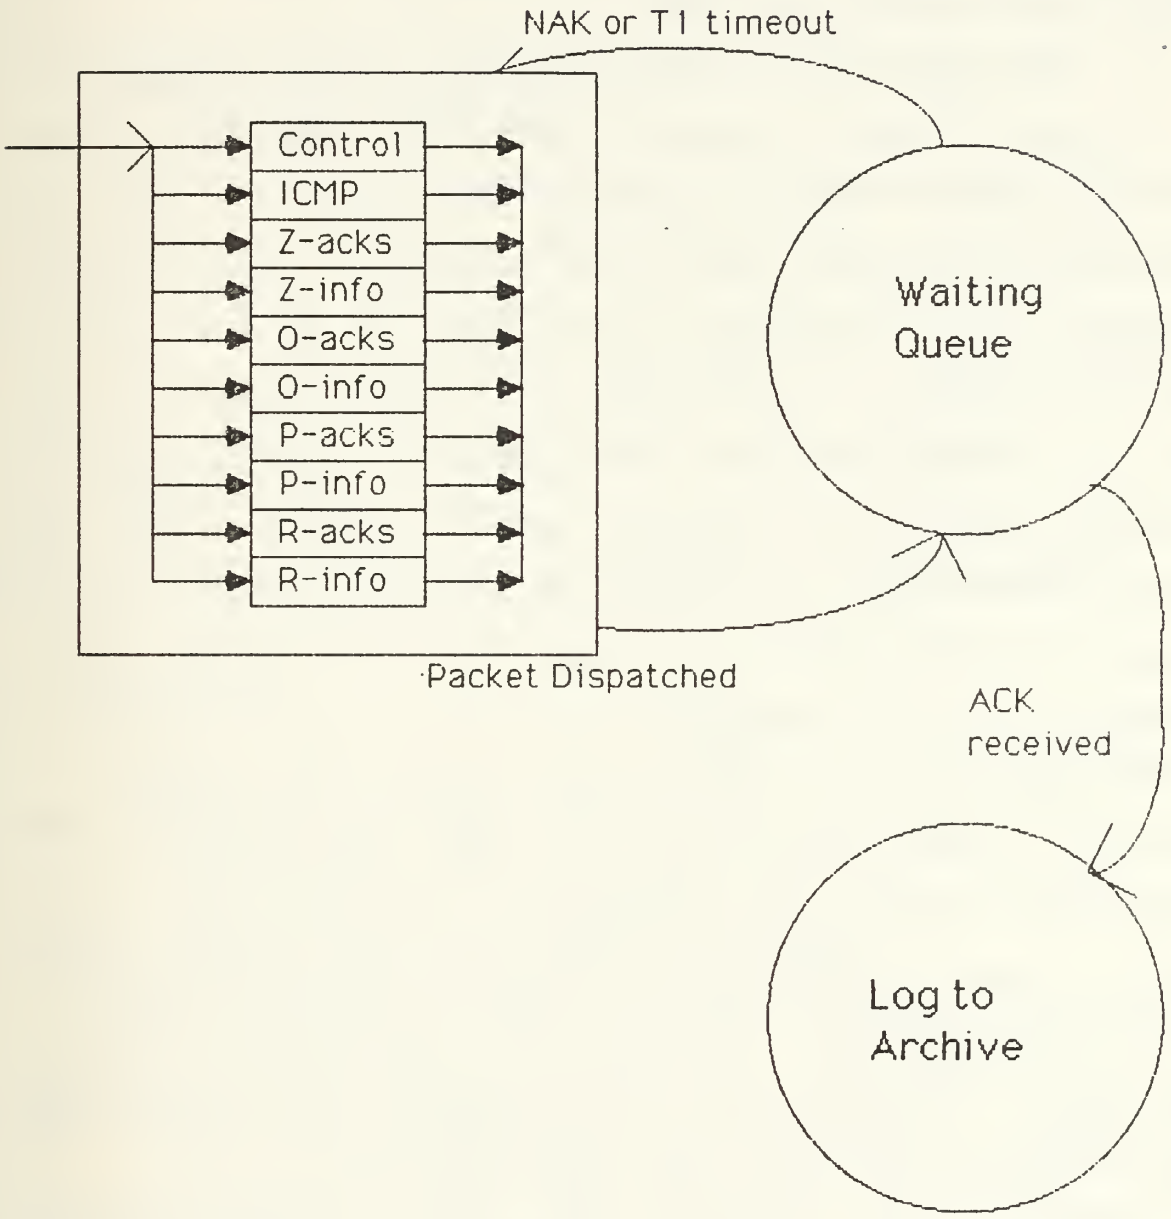

2. Counter

There will also need to be a counter as part of this structure. Each retransmission causes the counter to be incremented. At some point, the system either gives up, resets T1 to another value or informs the operator that communications with the receiver has been lost.

3. An Example of T1 and the Counter Working Together

A flash precedence message is sent. T1 is set quite short due to the grade of service required. The receiver is unable to acknowledge in a timely fashion (due to EMCON or simply speed of turnaround due to congestion). So T1 times out and the packet is sent again. After 3 transmissions, the counter causes T1 to be reset to a longer time (hours vice seconds). This corresponds to the current practice of rerunning broadcast traffic an hour after initial transmission. If at any time during this process, an ACK is received by the sender, the process is halted and the packet is logged as sent and received.

The receiver may or may not have received the packets correctly. If the counter runs to at least three before resetting T1, the receiver will have the material--three packets--to vote flawed packets (voting is explained in greater detail in Chapter Five). If at any point the receiver gains a valid packet, it originates an acknowledgement which causes the whole process to stop when the sender receives it.

If T1 pends (has not timed out) and the receiver is able to return a NAK to the sender, this causes immediate removal from the waiting queue and requeueing for transmission.

Parts of this algorithm can be omitted for NAK_only and No_ack packets; the full_ARQ algorithm is the most complete. The only caution is that NAK-only packets may have been logged to storage, only to have a belated NAK appear (due, perhaps, to the submarine finally being able to expose an antenna and transmit). The NAKed packets must be able to be retrieved and resent.

4. Multiple Addressees

This T1-counter structure may be multiplied for each packet. When collective addressing is used the sender must maintain a T1-counter structure for each addressee within the AIG or CAD.

The procedure is that a sender sends a packet marked for a collective address and then waits for each addressee to acknowledge in turn. Until all have acknowledged, the packet remains in the waiting queue subject to timeouts and NAKs.

F. CONCLUSION

Network level acknowledgement requires a Network Protocol that unifies several diverse links into a complete network. This chapter has indicated the requirements for such a protocol at the Network (lower half) layer of the ISO model.

This construction allows us to divide the communications system into two conceptual parts, links and nodes. Everything from the Network Layer (Network and Internet Protocols) up falls into the category of node. Nodal equipment, both hardware and software should be standard throughout the network. This means one-time development costs.

Links and Nodes

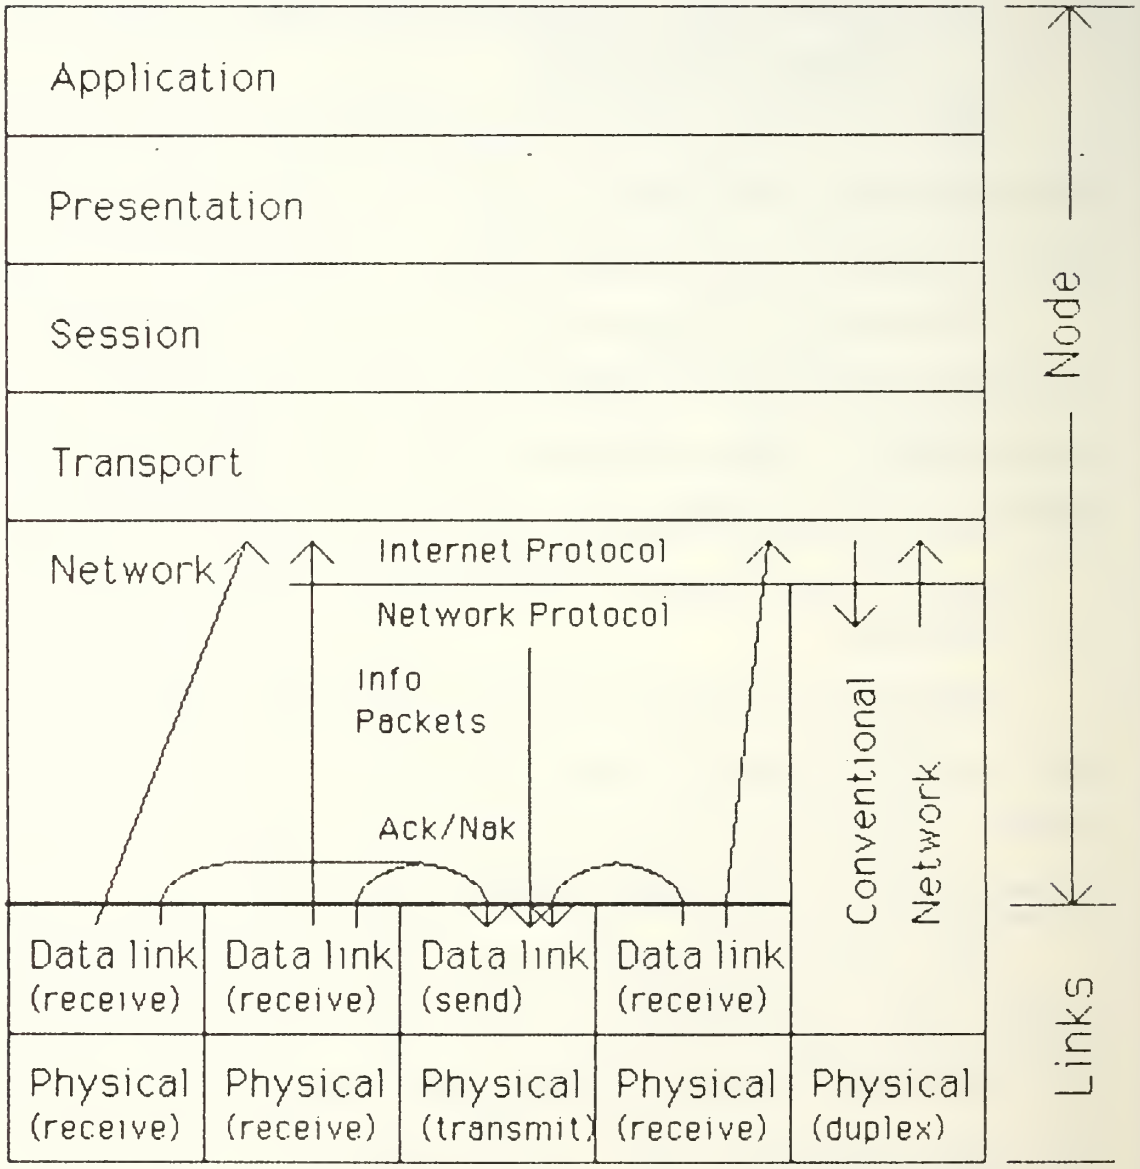

Logical link and physical layer hardware and software are unique to each link and can be customized to that link as necessary. Stretching the point a bit, even smoke signals could make an adequate link--one way--providing packets, both data and control, are passed to the Network Protocol.

Note that the existing transport layer - Internet Protocol - conventional LAN configuration is shown on the right side of the illustration. This remains as it currently exists--TCP/IP with a X.25 local network is a practical example. The point is that use of the Network Protocol to integrate heterogeneous one way links into a network is fully compatible with existing systems.

This allows an easy logical extension. Consider the illustration as a shipboard communications system. Radio central contains the node installation and the external links, such as HF, VLF and satellite, integrated through use of the Network Protocol. For intra-ship communication, an X.25 LAN can be used to connect terminals at various parts of the ship to the node at radio. A second X.25 network might be a practical way to serve a ship while it is in port.

Following chapters will concentrate on the logical link layers and the necessary protocol adjustments required at that level.

Network Access
Upward multiplexing

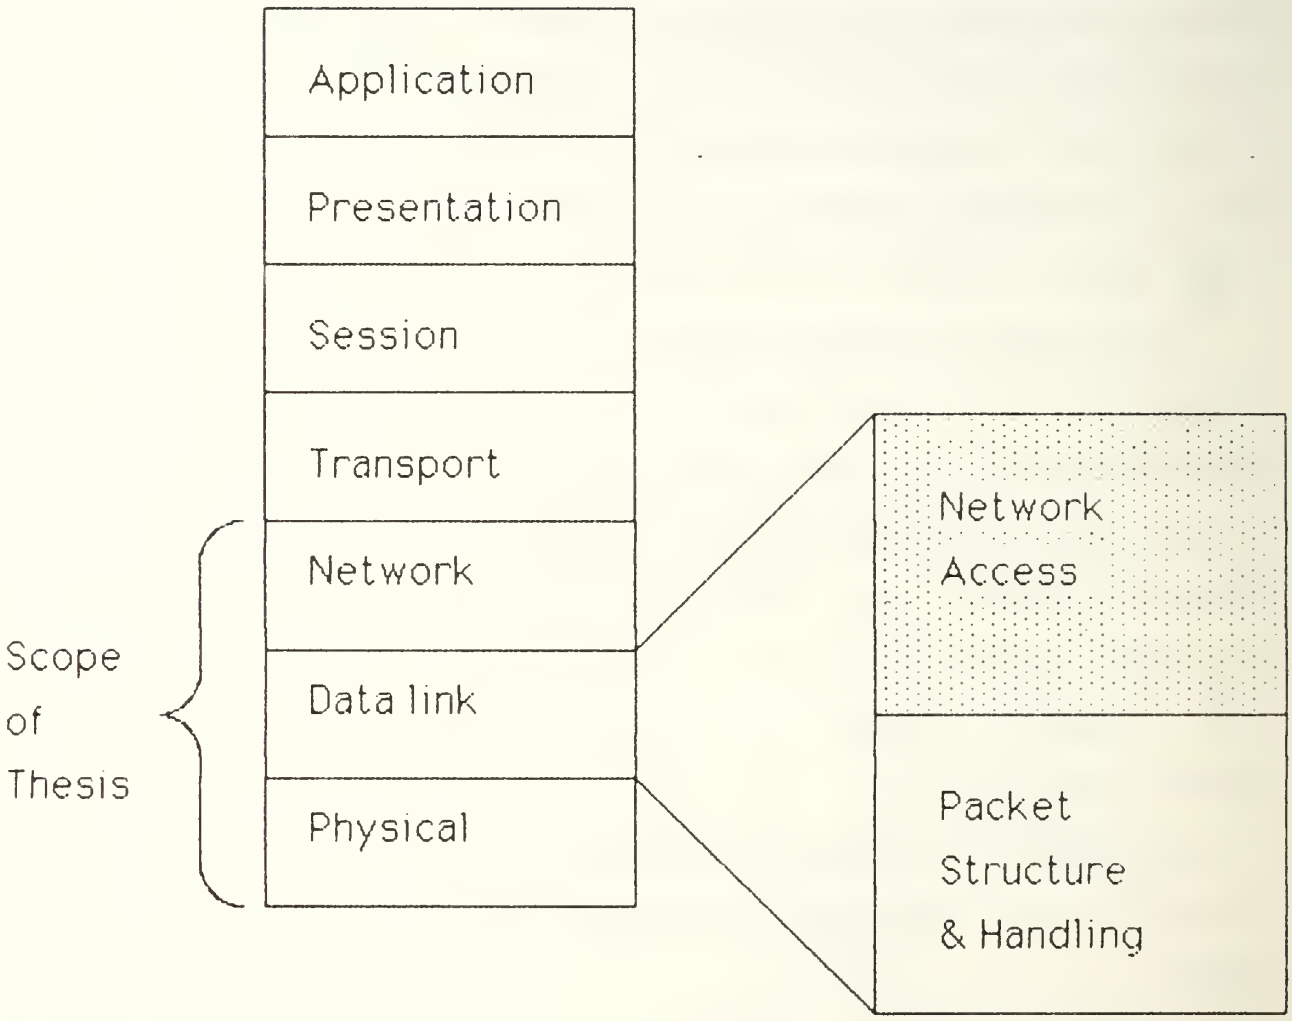

IV. NETWORK ACCESS--UPWARD MULTIPLEXING

A. DEFINITIONS

Upward multiplexing deals with several high level connections using a single, low level connection. Several ships and maybe the communications station will share a single frequency. How they share--who in a network uses a communications channel when--channel access--is the problem.

In any network, only one subscriber can transmit at any point in time. If this rule is violated, collisions occur and the transmissions are not successful. Network access protocols are designed to deal with this problem and prescribe rules to either prevent or deal with these collisions.

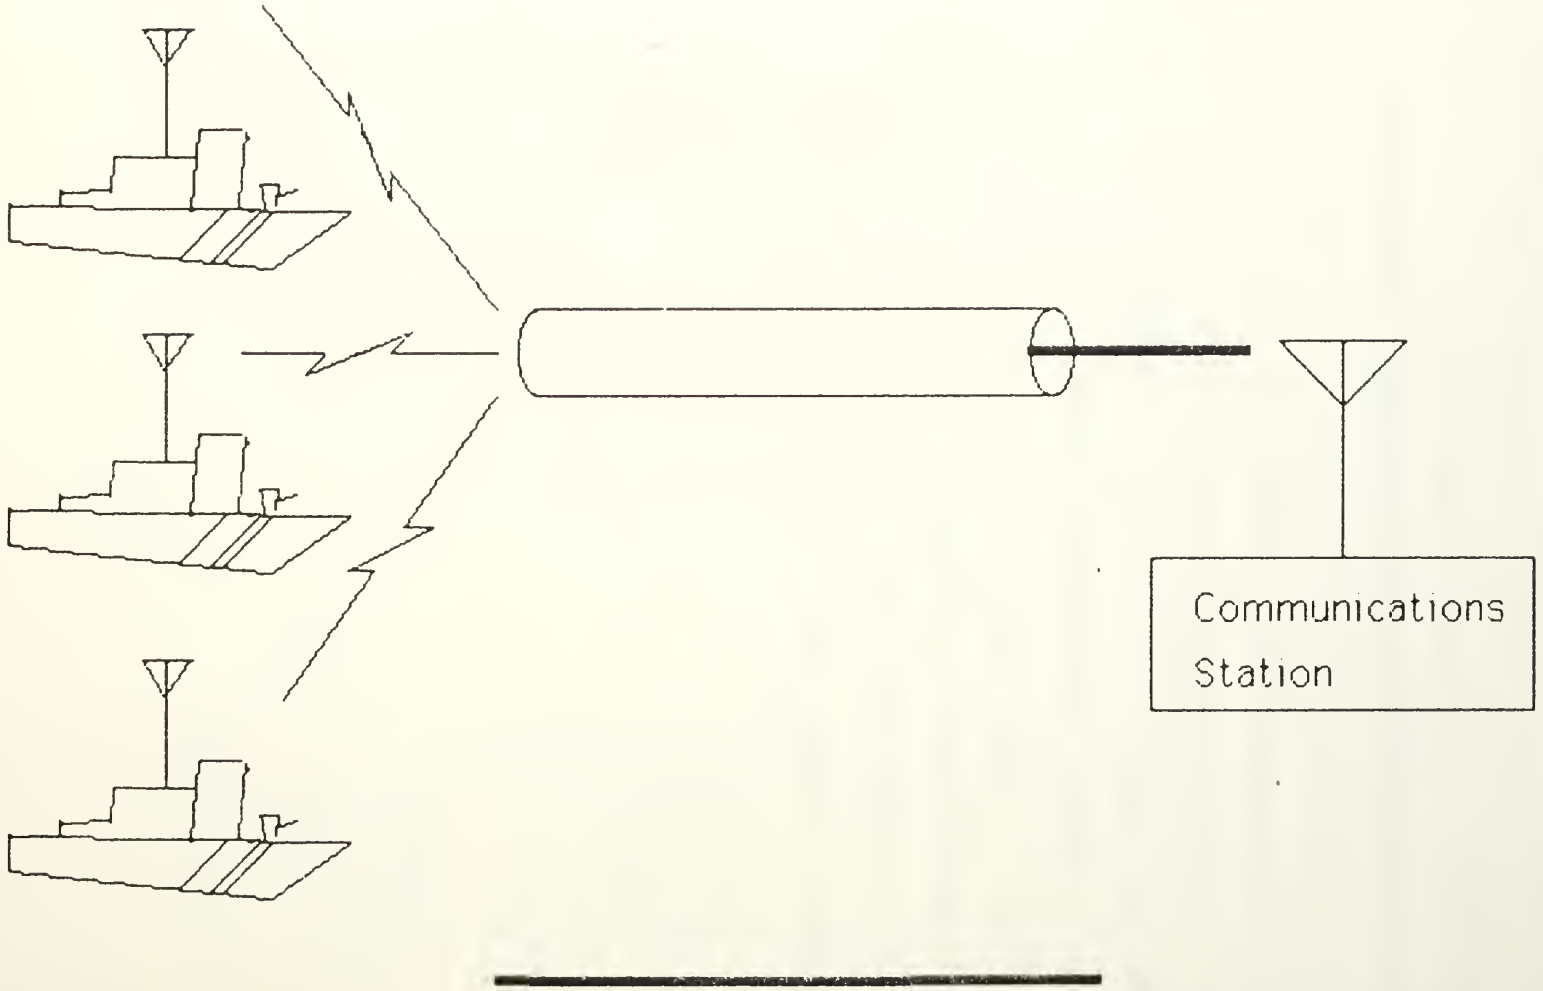

Upward Multiplexing

1. Scope

In this chapter, we will confine ourselves to the problems of HF communications. Because of the one way nature of the bands below HF, network access is not a problem--the communications station is the only transmitter on the network.

In the bands above HF, specifically satellite channels, a similar problem exist--stations cannot hear each other. Satellite communications can operate on some assumptions that do not hold true in the HF band, so we are dealing with the harder of the two problems. In any event, one solution is borrowed from the satellite system.

2. Flow Control

Since the reader should be convinced by the discussion in the previous chapters of the utility of network level acknowledgement, it will be assumed in this chapter.

Temporal linkage is an implication that follows network level acknowledgement. It is important in the consideration of the flow control mechanisms.

a. Breaking the Temporal Coupling

Conventional networks demand prompt responses. This is another characteristic that is built on the assumption of full duplex physical links. In a practical sea service environment, this may not be practical in many situations. Our objective is to make our system a tool of the tactician rather than making the tactician a slave of the communications system. Therefore the response time should be a function of the urgency and content of the communications itself, not an exigency of the communications system.

b. Flow Control Mechanisms

This means that Stop & Wait as well as sliding window protocols (Go Back N) no longer work. It will not do for a ship to be able to receive one (in the case of Stop & Wait) or eight (in the case of Go Back N as realized in X.25) packets and then have the system stop until the ship acknowledges those packets. The only feasible approach is Selective Repeat. This requires the more sophisticated packet identification scheme illustrated in our Network Protocol.

B. THE PROBLEMS

1. Conventional etwork ccess ethods nvalid

Stations in a skywave, long haul, HF network cannot be assumed to be able to hear each other. We can assume that they can all hear the communications station (at least most of the time). This invalidates the use of carrier sense multiple access (CSMA) schemes to manage access to the network.

a. Collision Detection

A more primitive network access scheme than carrier sense is collision detection (Aloha). By itself, collision detection is too wasteful of an already severely constricted bandwidth to warrant consideration as our primary means of controlling access.

Additionally, the instability of collision detection under an offered load greater than the channel capacity is a fatal flaw for sea service communications. This instability is described in several texts [Stallings, 85; Tanenbaum, 81].

b. Token Systems

Again, because ships may not be able to hear each other, token passing is not a viable network access strategy either. If the token is passed via the communications station, the system becomes essentially a centrally controlled one.

2. Physical Layer Considerations

As we discussed in the last chapter, simulating full duplex is difficult to impossible in sea service HF communication links. Additionally, within the HF band, there are some further considerations.

a. Synchronous Equipment

The synchronization preambles required by various equipments make it expensive to reverse the channel (switch from receive to send). Most conventional networks operate in an asynchronous environment where this reversal is quick and painless.

But once a network becomes a covered circuit--link encrypted--synchronous communications are imposed by the KG-84 cryptographic device. Additionally, there is a power-up time required by high powered transmitters, plus a synchronization requirement for synchronous modems. These problems are theoretically possible to deal with at the physical level, but the solutions are more expensive than dealing with them at the logical link and network levels.

The modems and cryptographic equipment most likely to be used are of the type that require synchronization. For instance, the USQ-83 modem and the KG-84 cryptographic device each require 0.8 seconds of synchronization or 1.6 seconds for each

transmission sequence. One 256 byte packet transmitted at 2400 baud requires 0.85 seconds. Thus this (purely arbitrary) choice of parameters would result in nearly twice as much channel time taken up in synchronization than in transmitting data.

This situation can be made considerably worse by selection of modems. Some modems offer deep interleaving, or spreading of a message over a longer period of time, to combat fading errors. While this approach improves the quality of a continuous bit stream, it can potentially severely degrade a logically bursty communications system.

For this reason, we wish to construct the overall system so that at the physical layer reversals are minimized. That is, the physical layer bit stream is as continuous as practical, regardless of the logical 'burstiness' at higher levels.

A practical solution is to bunch several packets together. Indeed the X.25 protocol allows the ending flag byte of one packet to serve as the leading flag byte of the next. A series of flag bytes can serve as spacers between packets so the physical layer equipment does not lose synchronization. This is the first step toward making a logically bursty circuit appear continuous at the physical level.

This synchronization consideration is an argument for long access periods for each user--long enough to clear all traffic in the queue. The opposing consideration is the requirement to allow access to all users--one user should not be able to monopolize the channel. This is the tradeoff between throughput

and grade of service. The requirement to handle high precedence traffic requires a degree of interruptability.

b. Full Duplex Inhibition

We will continue to observe the principle that ships will not be able to transmit and receive at the same time. This requirement has several sources:

(1) Signal to Noise. The physical difficulties associated with receiving a faint signal in the presence of a strong transmit signal. This has been discussed earlier.

(2) EMCON. The inhibitions against transmitting that may be imposed upon a ship by tactical emission control (EMCON) considerations.

(3) Equipment. Physical layer equipment--the radios themselves--are often packaged as tranceivers. With both the transmitter and receiver in the same physical box, it becomes impossible to operate both simultaneously.

(4) Others. Several other reasons such as hazards of electromagnetic radiation to ordnance (HERO), incompatability with other ship sensors, and shipboard power failures have the same effect--restriction on transmission.

c. Controlling Downward Multiplexing

If a ship needs more capacity than is available on one channel, then multiple channels must be 'ganged together' to provide the necessary bandwidth in aggregate. In the HF band, one controller must control all of the shipboard transmitters to preclude one channel transmitting while another is attempting to receive.

Since this synchronization issue is a link one specific to HF, it should be handled at this level rather than at the network level where the protocol should not be required to know about dependencies among links.

The remainder of this chapter deals with this problem by building a model flexible and robust enough to use in our HF system.

The first, and most fully developed one builds from the CUDIXS network access algorithm. A second alternative is derived from the Link 11 method.

C. A SIMPLEX POLLED CIRCLE

1. Cycle Description

All stations (including the communications station) operate on the same frequency, both send and receive.

a. Step 1) Network Poll

The communications station sends a polling packet to each station in turn. The ship station responds with a summary of its queue contents. For example, 250 bytes of Immediate, 1000 bytes of Priority and 2500 bytes of Routine.

b. Step 2) Sequence Order List (SOL) Organization

The communications station balances the ships' queues with its own outgoing traffic and prepares a schedule. This is quite analogous to the operations of a scheduler in a priority driven multi-processing operating system--the scheduler is analyzing the ready list. A different, but equally valid, analogy is that the prioritization algorithm is simply a mechanization of the radioman's rules that are used today. The

communications station then broadcasts the SOL which gives each ship her turn.

c. Step 3) Ships Send Traffic in Turn

The cycle is time-clocked so ships are not required to be able to hear each other. Acknowledgements for the previous cycle are included here. Note that timing requirements are on the order of milliseconds--easily accomplished by the processors, but not requiring precise time.

d. Step 4) Communications Station Turn

The communications station sends its traffic along with acknowledgements for ship station traffic. Requests for repeated packets (NAKs) can be honored at this time.

e. Step 5) Iteration

The cycle is finished and the next is started by the communications station proceeding with Step 1).

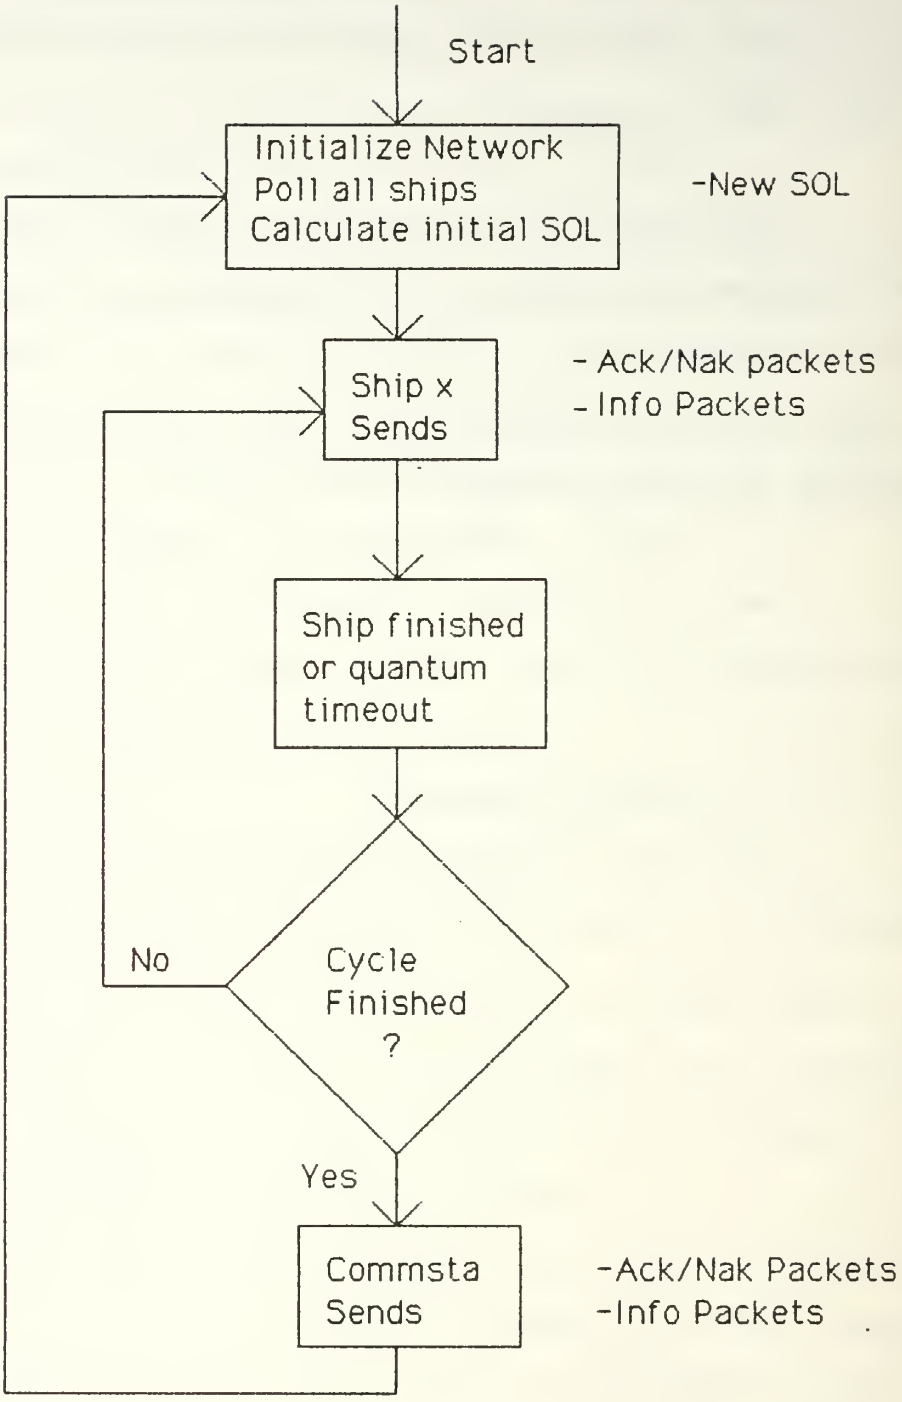

Basic Simplex Polling

2. Sensitivity Analysis

This results in a requirement for each ship to transmit its synchronize sequence twice per cycle. The communications station will have to transmit its synch once per ship in the polling sequence plus twice per cycle (once for SOL transmission and once for its traffic). This results in about 85 seconds of the two minutes gobbled up by synchronization requirements. This is about 3/4 of the cycle--obviously not good enough. This basic model is presented as a place to start.

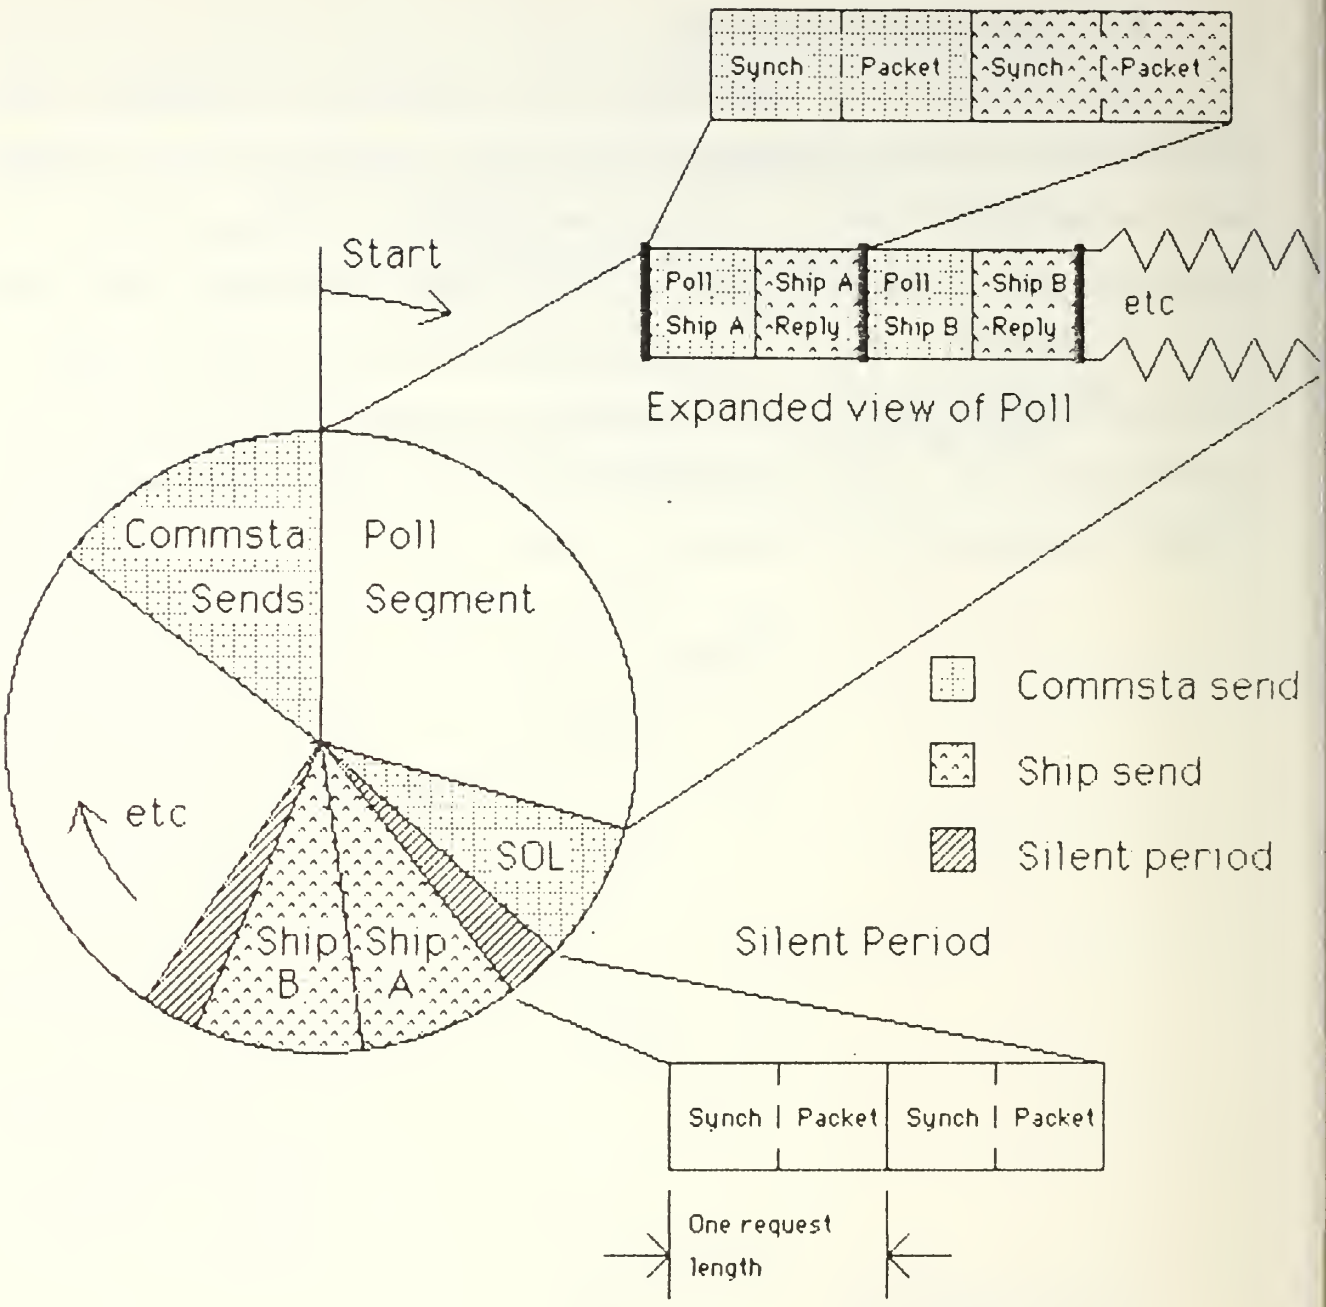

Basic Simplex Polling circle

D. A REVISED SIMPLEX POLLING SYSTEM

1. Revised Cycle Description

This is a simplex description that eliminates some of the overhead due to synchronization.

a. Step 1) Network Poll

The communications station sends a polling packet to each station in turn. Each ship station responds.

b. Step 2) Sequence Order List (SOL) Organization

The communications station balances the ships' queues with its own outgoing traffic and broadcasts a schedule--the SOL--which gives each ship its turn.

c. Step 3) Communications Station Transmit

Communications station sends its traffic along with acknowledgements for ship station traffic. By putting SOL broadcast and communications station traffic together, the requirement to synchronize a second time is eliminated. The communications station would acknowledge receipts from the previous cycle at this time.

d. Step 4) Ships Send Traffic in Turn

Acknowledgements for the previous cycle are included here. At the end of its turn, each ship station includes a polling response packet indicating the amount of traffic remaining in that ship's queue.

e. Step 5) Iterate

The cycle is finished and the next is started by the communications station proceeding with Step 2). The polling cycle is not required because the data was gathered in Step 3).

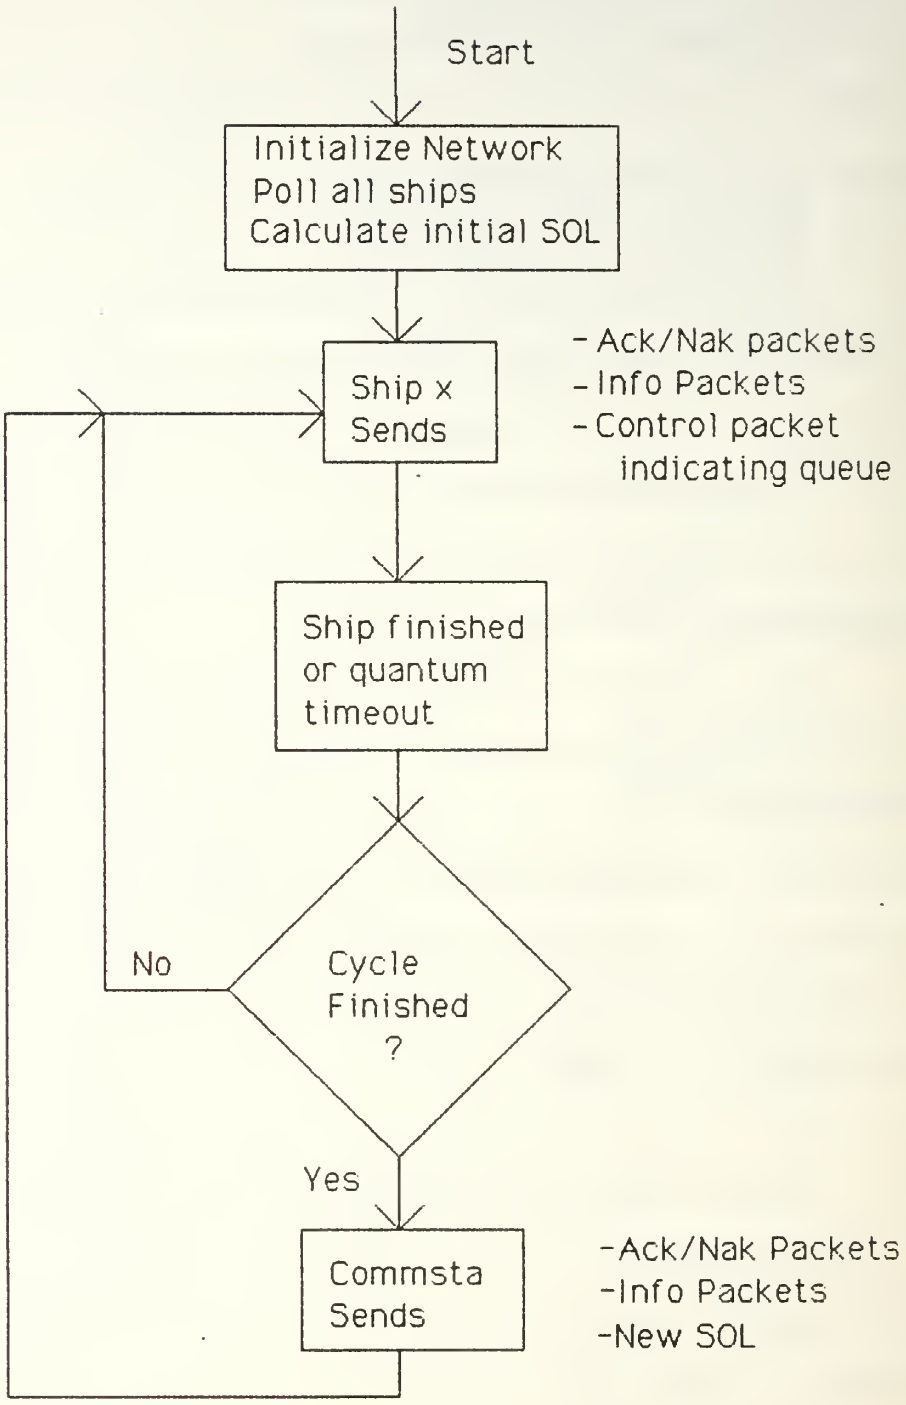

Revised Simplex Polling

2. Sensitivity Analysis

This will result in a slight time lag in the queue management as data that is added in the interim will not have been accounted for in the poll. However this problem is not serious; if the ship has managed to send all of the Immediate and Priority packets but has several Routine packets remaining, then Immediate and Priority packets added to the queue in the interim would be sorted to the top of the queue and be first out in the next cycle anyway. If two or three cycles elapse before a routine packet succeeds in clearing the queue and getting transmitted, the delay is inconsequential.

The slight sacrifice in grade of service is more than made up in increased throughput which, if it eases congestion, restores the grade of service to better than the level expected under the basic algorithm.

In this cycle, each ship must send synch once and the communications station once. Synch time is reduced to 30 seconds or about 26% of a 2 minute cycle for a hypothetical 20 ship network.

If communications can proceed at a full 2400 baud during the remaining 90 seconds of the cycle, approximately 27k bytes could be carried over the circuit. One page of text, a naval message if you wish, is around 2k bytes in length meaning that approximately 15 one page messages could be sent over the link in the two minute cycle. (See the summary chart at the end of this chapter for estimations for each technique.)

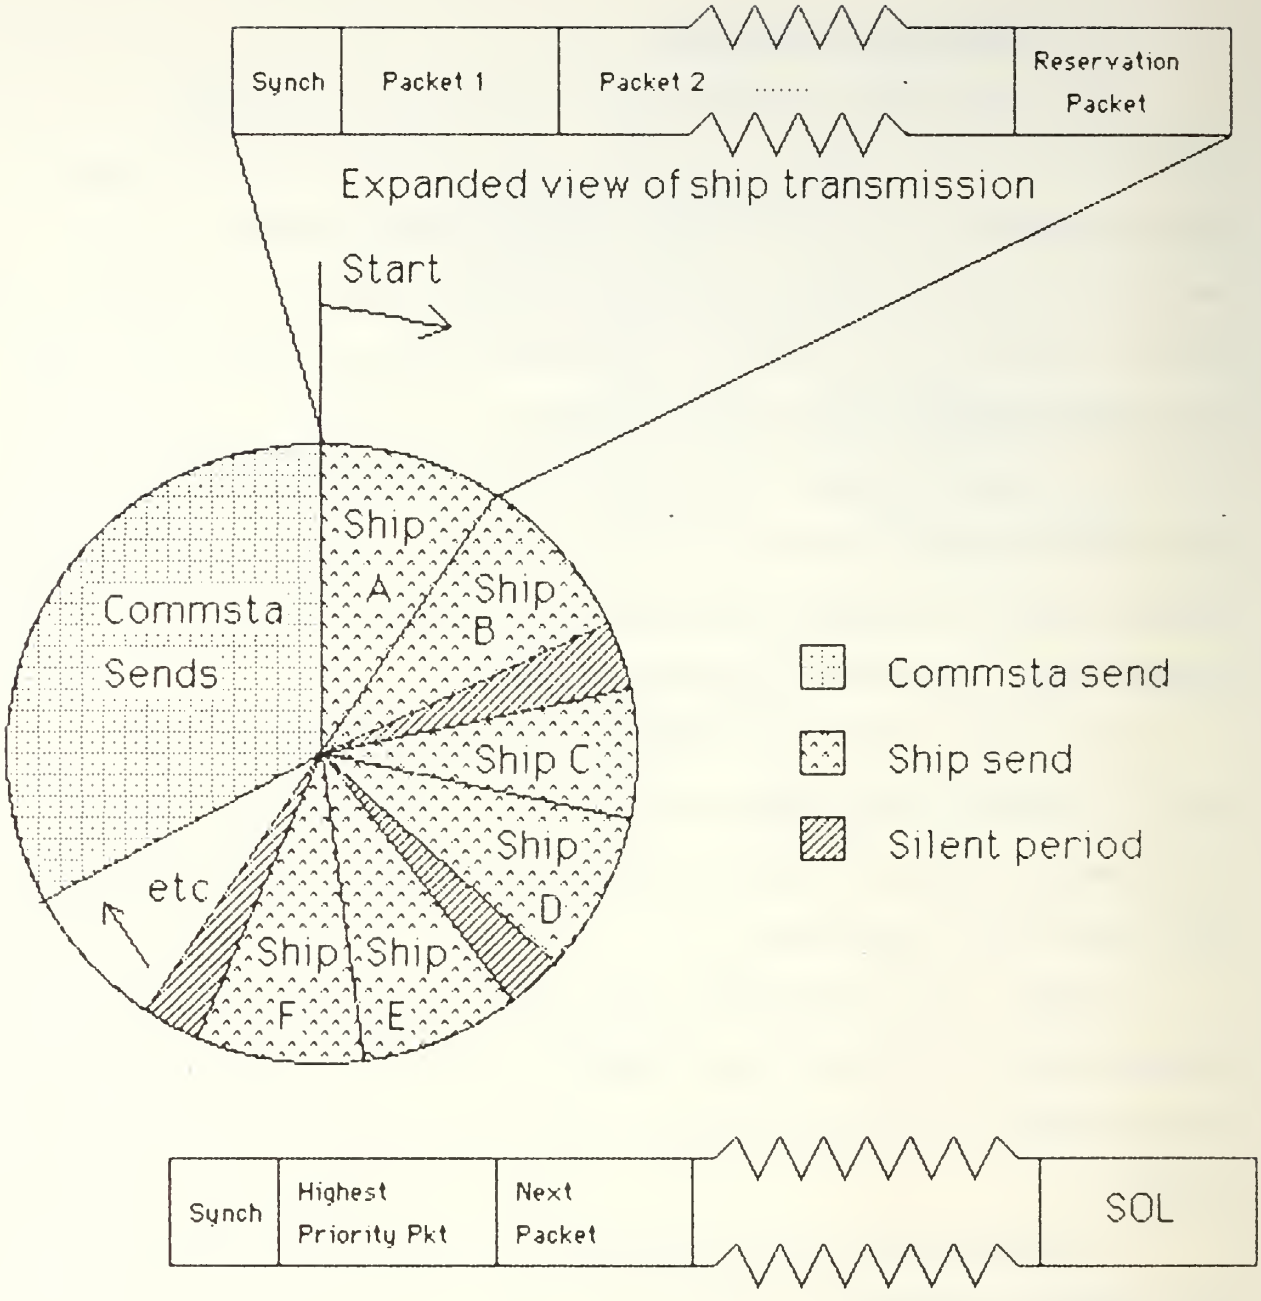

Revised Simplex Polling circle

3. Net Entry--A Second Tuning

A ship not in this cycle that wishes to enter it cannot find an opening in the above described cycle. So the basic cycle is modified slightly. When the SOL is broadcast by the communications station, provision is made for some silent periods where none of the current participants in the net are transmitting. Silent periods are a standard part of CW communications on calling and distress frequencies where certain periods each hour are set aside for listening for faint signals and are not to be used for calling. These silent periods are called random access time slots (RATS) in the CUDIXS system.

The silent periods exist for two purposes:

1. a new user sends a net entry packet to the communications station during this period. Since more than one new user may wish to enter the net, a collision detection system must be used. If a collision avoidance system designed to detect the synch tones of a competing ship station were also included, silent period use would be more efficient. The net entry packet contains the same information that is polled for by the communications station: contents of the traffic queue.

2. any user can get off a brief (one packet) high precedence message. This allows any station, ship or communications station, to transmit a Flash packet without waiting its turn.

It should be pointed out that the cycle would normally be full on a fully loaded network--the offered traffic is greater than the capacity of the channel. If the channel is not fully loaded, the communications station simply allocates the idle time to larger silent periods. In this way, a decreased traffic load will result in improved grade of service because the silent periods are available to any ship on an immediate basis.

4. Balancing Between Trade of Service and Throughput

The more silent periods in a cycle, the less the throughput. But, more silent periods mean that extremely high precedence traffic should get through with very little queue delay. Thus a ship-communications station network might be best balanced for throughput with small, infrequent silent periods. By contrast, a network serving aircraft or submarines may not need high throughput, but rapid access to the net is imperative. Here, large, frequent silent periods and short scheduled transmission periods would be better. This balancing should be performed by the operators at the communications station, not the system engineers at design time.

Similarly, if a network is to be used for large quantities of data transfer, silent periods might be omitted entirely. Network access would be gained by a ship sending an entry request to the communications station by another means--a calling channel perhaps--and then copying the frequency in question.

Balancing the tradeoffs between grade of service and throughput can also be done by varying the cycle time. A shorter cycle time improves grade of service at the expense of throughput. The number of synchronizations per cycle remains constant, but the number of cycles in a period of time increases.

5. Full Period Terminations

A full period termination is an extreme version of the cycle where a particular channel is dedicated to serving one ship. In this case, no silent periods are needed. The communications station and the ship exchange queue status packets along with their traffic. If the channel is turned around at least once per

minute, no intolerable delays are encountered, and each station must synch once per cycle with negligible overhead. The even more extreme case where one user requires multiple channels for adequate capacity takes us to the downward multiplexing situation which was the subject of previous chapters.

6. Missed Schedules

A ship that fails to receive the SOL must refrain from transmitting in the cycle. Conversely, if a ship misses a scheduled transmission, the communications station must allocate it some time in the succeeding cycle. If a ship continues to miss transmissions, it may have entered EMCON unexpectedly or suffered a power failure or some other casualty. After a polite number of cycles, the communications station can delete its time slot in the SOL in favor of other users. However, because the ship may be receiving passively, the communications station should continue to send packets to the silent ship until it is advised otherwise.

7. Customized Service

By use of control packets, the ship may arrange several variations of service:

1. The ship may enter an EMCON period that is forecast. Therefore, it will wish to receive, but can signal the communications station not to include it in the SOL until the ship reenters the net via a silent period.

2. The ship may anticipate traffic and request an SOL time slice every cycle even though it does not have pending traffic. This spares the ship from the requirement to reenter the net when the traffic is ready to send.

3. The ship may include a control packet on one network initiating a login into a new network. This might occur in anticipation of a large volume of traffic requiring use of more than one channel. It might also occur in anticipation of a diurnal frequency shift which will render the existing frequency unusable at a future time.

8. Advantages

This second iteration of a simplex protocol is entirely practicable and is reasonably efficient. It gets reasonable amounts of throughput through a single HF channel. The duplexing problems of small ships and transceivers are avoided.

9. Disadvantages

All ships in the network are tied to the same frequency. This problem is academically intractable, but practically inconsequential as several frequencies--thus several available links--could be expected to be in use at any one time. A ship needing to change frequencies simply logs out of one network, retunes radios, and logs into a new network.

E. A FULL DUPLEX SYSTEM

In a duplex system, the communications station would poll each station for traffic load pending, as before. Then, rather than publishing the SOL, it simply converses with each ship, one at a time, in turn. This conversation would be full duplex--the subject ship is responding as the communications station is sending. Both stations attempt to clear their traffic queues for the other during this window. At expiration, the communications station goes on to the next ship. Any unsent traffic remains in the queue for the next cycle. Any sent but unacknowledged packets are presumed lost at sea and are re-queued in the next cycle.

This system has some real attractions to it:

1. A standard X.25 protocol would work for this system. Similarly, a Go-Back-N error control algorithm with its modulo arithmetic sliding window remains useful, if inefficient.

2. Ship stations needn't be terribly sophisticated regarding timing and network access. Each simply waits its turn until called upon by the communications station to send its traffic. In particular, except for finding the silent periods, there are no shipboard timing considerations.

3. Probably the biggest plus to this system would be that there is no requirement for the ships to be on the same frequency. Assuming that the communications station has the necessary frequency agility in its transmitters and receivers, it can shift to a new frequency for each ship if necessary.

As a practical matter where the Navy usually sails ships together, this advantage may not be particularly useful. Ironically, for Coast Guard use, this is the biggest potential advantage.

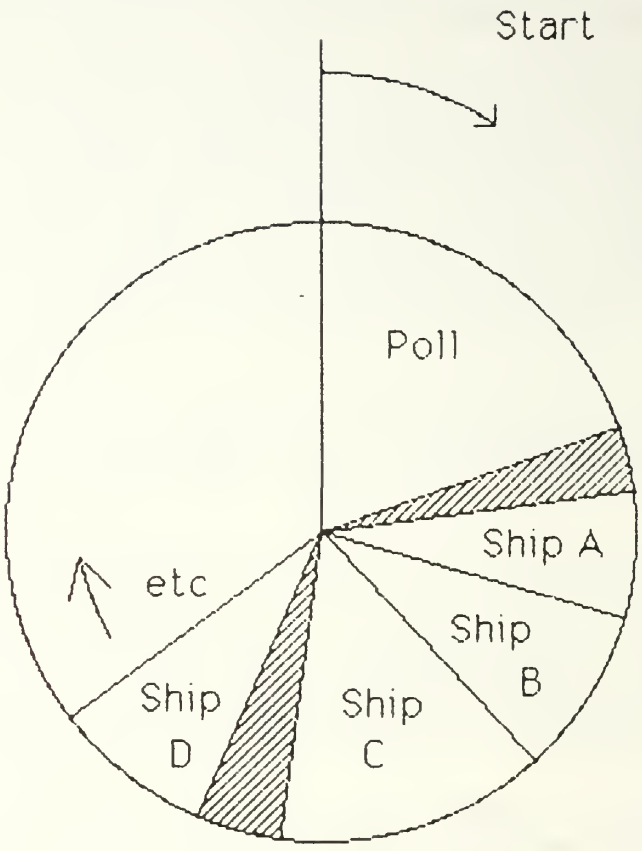

Full Duplex Polling circle

Duplexing has some fatal drawbacks, however. The worst is that it will not work. We have discussed previously the inability of most ships to operate full duplex in the HF band.

The second drawback is the close temporal coupling implied between sending and receiving. This makes the downward multiplexing described in the last chapter impossible. In particular, there would be no way to integrate the fleet broadcasts into the system.

The third drawback is that a ship would be unable to receive passively. It is not tenable that a ship be required to broadcast in order to receive.

Collective broadcasting will not work in this model as a packet received by an addressed ship cannot be acknowledged if the ship is not currently being worked by the communications station. For sea service communications, this represents a serious inefficiency.

Full duplex turns out to be both less efficient and less flexible than the half duplex model developed below.

F. HALF DUPLEX MODEL

1. Cycle Description

This is an expansion of the simplex description that combines the advantages and avoids the fatal drawbacks of full duplex.

In particular, note the loose synchronization between the send and receive channels. It now becomes possible for one of the channels (presumably the ships' send channel for ESM reasons) to be a satellite channel while the communications station continues to use an HF channel. Once a half duplex arrangement is in place, the multiple channels of the downward multiplexing concept of earlier chapters is also fully usable.

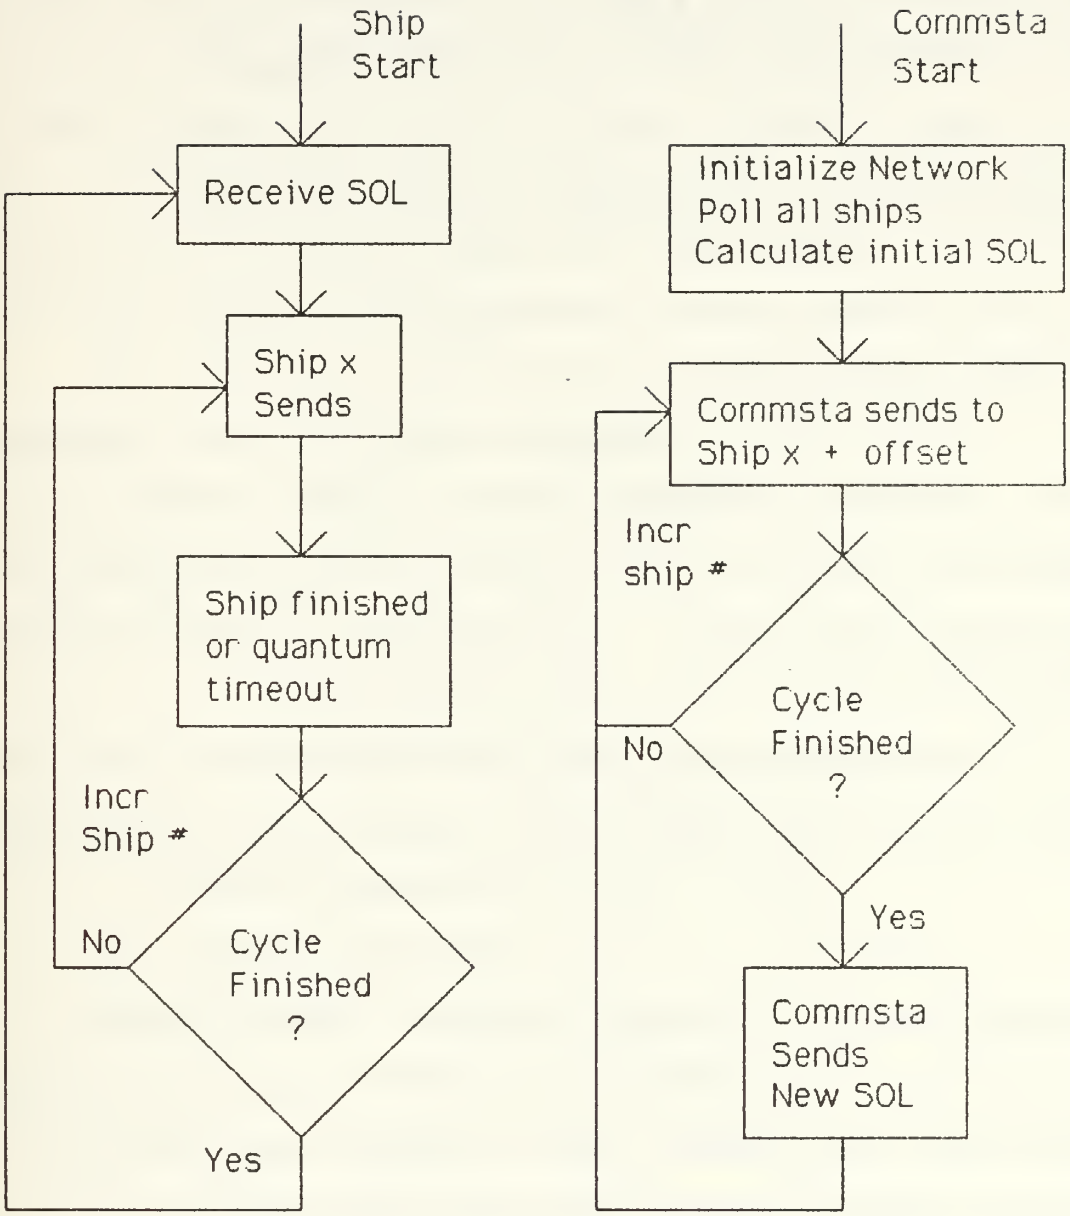

Half Duplex Polling

a. Step 1) Network Poll

In order to initialize and start a network, the communications station sends a polling packet to each station in turn. The difference here is that the communications station is using a send frequency where it is the only talker. Each ship station responds with a summary of its queue contents, as before. One way to accomplish this is simply to announce the network and allow any users to log in to it--initially, the circuit will be totally silent period.

b. Step 2) Sequence Order List (SOL) Organization

The communications station needn't balance the ships' queues with its own outgoing traffic any more, as they are not sharing the same channel. The communications station broadcasts the SOL which again gives each ship its turn.

The primary difference of a half-duplex cycle comes here. The next two steps are not consecutive, but concurrent.

c. Step 3) Ships Send Traffic in Turn

The cycle is time-clocked, as before, so ships are not required to be able to hear each other. Acknowledgements for the previous cycle are included here. At the end of its quantum (turn), each ship station sends a control packet indicating the amount of traffic remaining in that ship's queue.

a) There are four options for a ship station at this point. The first is that the ship has traffic remaining in the queue and needs space in the next cycle.

b) The second option is that the queue is exhausted, but the ship wishes to remain in the cycle because it may have traffic ready to send by the next cycle and wishes to avoid the time required to reenter the net.

c) The ship wishes go silent but remain in the net. When it has more traffic to send, it has a minimal space in the SOL to take up the thread. This option has the advantage of not requiring the ship to transmit at all, thus compatible with possible EMCON requirements. The ship can receive without an immediate requirement to respond.

d) The ship checks out of the net. To reenter, it uses the silent period procedure. Any shore-ship traffic is stored unsent by the communications station until the ship logs back in.

d. Step 4) Communications Station Transmit

The communications station sends its traffic along with acknowledgements for ship station traffic. But traffic to a ship station is timed so that it is not being sent at the same time that the ship is transmitting. As soon as more than two or three ships are on the net, this becomes fairly easy to do--if we assume that all the ships have roughly equal amounts of traffic, simply offsetting the communications station transmit packets by about half of the cycle time would prevent interference in all but the most distorted schedules.

e. Step 5) Iterate

The cycle is finished and the next is started by the communications station proceeding with Step 2). The polling cycle is not required because the data was gathered in Step 3). If the communications station broadcasts the SOL packet twice during a cycle, each ship should be silent during at least one instance, so each can receive it at least once.

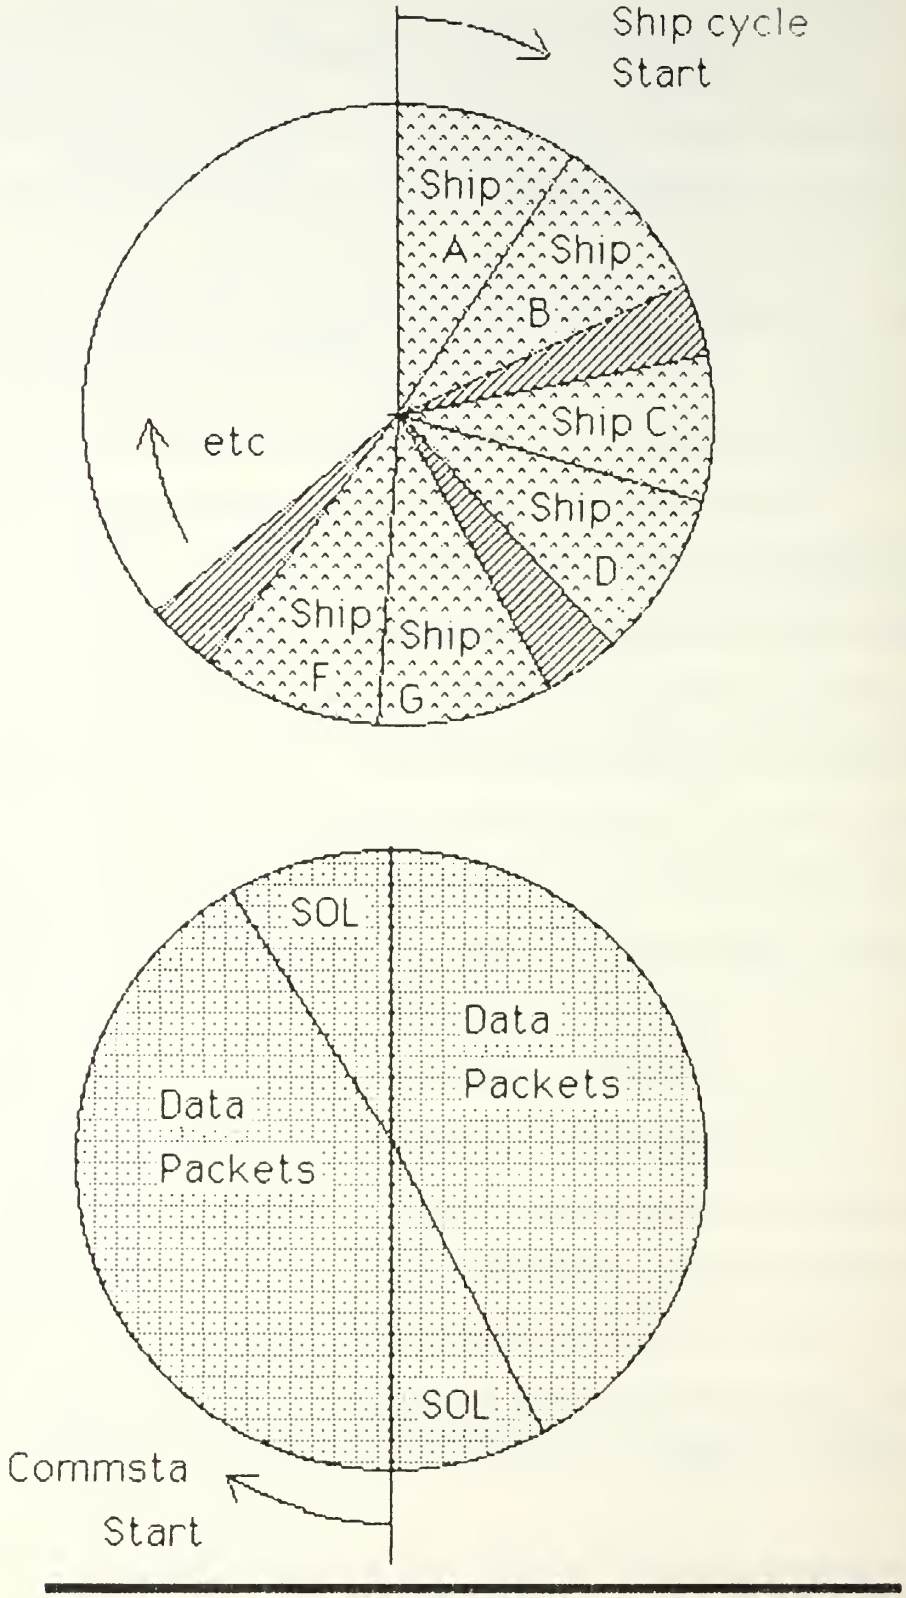

Half Duplex Polling Circle

2. Sensitivity

The communications station theoretically needn't synchronize at all during the cycle as it is broadcasting continuously. Practically, however, at least one synch per cycle may be necessary to accomodate new users and to recover ships that may have lost synchronization. This would depend on the requirements of the physical layer equipment.

Ship stations also only need to synchronize once per cycle. Thus the total synchronization overhead per cycle will be equivalent to the total number of stations in the network.

If communications can proceed at a full 2400 baud during the full 120 seconds of the cycle, approximately 36k bytes should be carried in the shore-ship direction and 27k bytes in the ship-shore direction over the circuit.

Given a one page text message of around 2k bytes in length, approximately 18 messages could be sent in each direction over the link in the two minute cycle.

3. Net Entry

A ship not in this cycle that wishes to enter it again uses the silent periods. Since the SOL is broadcast by the communications station using of a collective (CQ) call, any entering ship can determine when the silent periods occur and can send entry packets then.

4. Advantages

The most important is that throughput is brought near to the limit that the channel and physical layer equipment will allow. Half-duplex is somewhat more efficient than the simplex model. Additionally, a simplex configuration can be easily expanded by

the operator into a half-duplex and on to a half-multiplex one as capacity needs require.

The send/response cycle is partially decoupled. The only temporal coupling is the requirement for a ship to copy the SOL in order to transmit. Thus it is now practical for the additional multiplexing to be added, including integration of satellite channels into the system.

The small ship syndrome problems inhibiting full duplexing are avoided.

Voice conversations become fully practical. This would be done by making the entire cycle silent period with each converser entering (with the push-to-talk button perhaps) in a collision detection/avoidance mode. Using ISO model parlance, access control is delegated upward to the user level. Voice packets would be coded as No_ack ones so that the receivers ARQ is inactive.

If the noise level on a single channel is too high to allow adequate bit rates on a single channel, multiple channels can be ganged together with a portion of the voice packets being sent over each. The network layer at the destination merges the packets back together into a continuous stream that can then be directed to the synthesizer where voice tones are reconstructed.

It would still be possible to get data through the channel using the 'blank spots' in conversation, but only in a collision detection mode. Expeditious delivery or respectable throughput could not be promised on a voice circuit.

When unbalanced traffic loads occur, the number of channels in each direction can be adjusted to match the load. For instance, suppose that the notional 20 ships have more traffic than the one ship-shore channel allows. A second adjacent frequency can be utilized by moving half of the ships to that. Assuming that the shore-ship load remains small enough to fit in one channel, the only adjustment to that channel is that two SOLs, one for each ship-shore channel, must be sent vice the one for a simple half duplex system.

G. A CUED ACCESS MODEL

A network access model similar to that employed by Link 11 in the Naval Tactical Data System can also be practical. This system is undoubtedly practical, but requires some fast scheduling footwork by the communications station. This is especially so when several ship-shore links are being supported by one shore-ship one.

1. Polling and Initiation

After the network is organized, the communications station signals the first ship to commence transmissions.

2. Transmission

The ship sends until it has exhausted its queue or time is up. An end of transmission (EOT) packet brings up the rear. Included in the EOT packet is the queue status and other control information needed for the next cycle.

3. Signalling the Next Ship

When the communications station receives the EOT packet from the first ship (or concludes that quantum expiration has been

reached in the absence of reception from the ship), it signals the second ship to commence transmissions. If the network is configured in a simplex fashion, this entails a circuit reversal which requires initiation of a send synchronization by the communications station. More practically, the network would be configured in a half duplex arrangement allowing the communications station to transmit continuously. Upon receipt of the EOT, the communications station simply interrupts its queue of packets and inserts the necessary signal packet.

In the same manner, the communications station can signal a silent period to the network. At the conclusion of the silent period, the communications station would signal the next ship to start transmitting.

If a ship misses a transmission, it would be skipped and given another opportunity either in the next cycle or later in the current cycle.

This algorithm is quite practical for a half duplex HF arrangement. But it appears to require a closer temporal coupling between the two channels than the CUDIXS based system. It also requires minimal processing and propagation delay in order to be efficient. Processing delay should not be a problem, and propagation delay is not prohibitive in the HF band as it is in satellite systems.

Consequently, this appears to be an effective algorithm in the HF band, but somewhat less efficient when either channel involves satellite propagation.

Since it is not vital to pick between the two algorithms at this point, both are presented to the reader (and possibly to the system builder).

H. CONCLUSION

This chapter has presented three practical algorithms for network access control in an upward multiplexed environment. The tuned simplex and the half-duplex methods should be servicable in any band. The cued access method should work well in the HF band but will be less efficient than the current CUDIXS model used in satellite bands.

In particular, the half-duplex method should serve well in a fully integrated, network-layer acknowledged, communications system described in the previous chapter.

The concluding table shows comparative throughput for each of the algorithms presented.

Comparative Throughput for Different Access Schemes

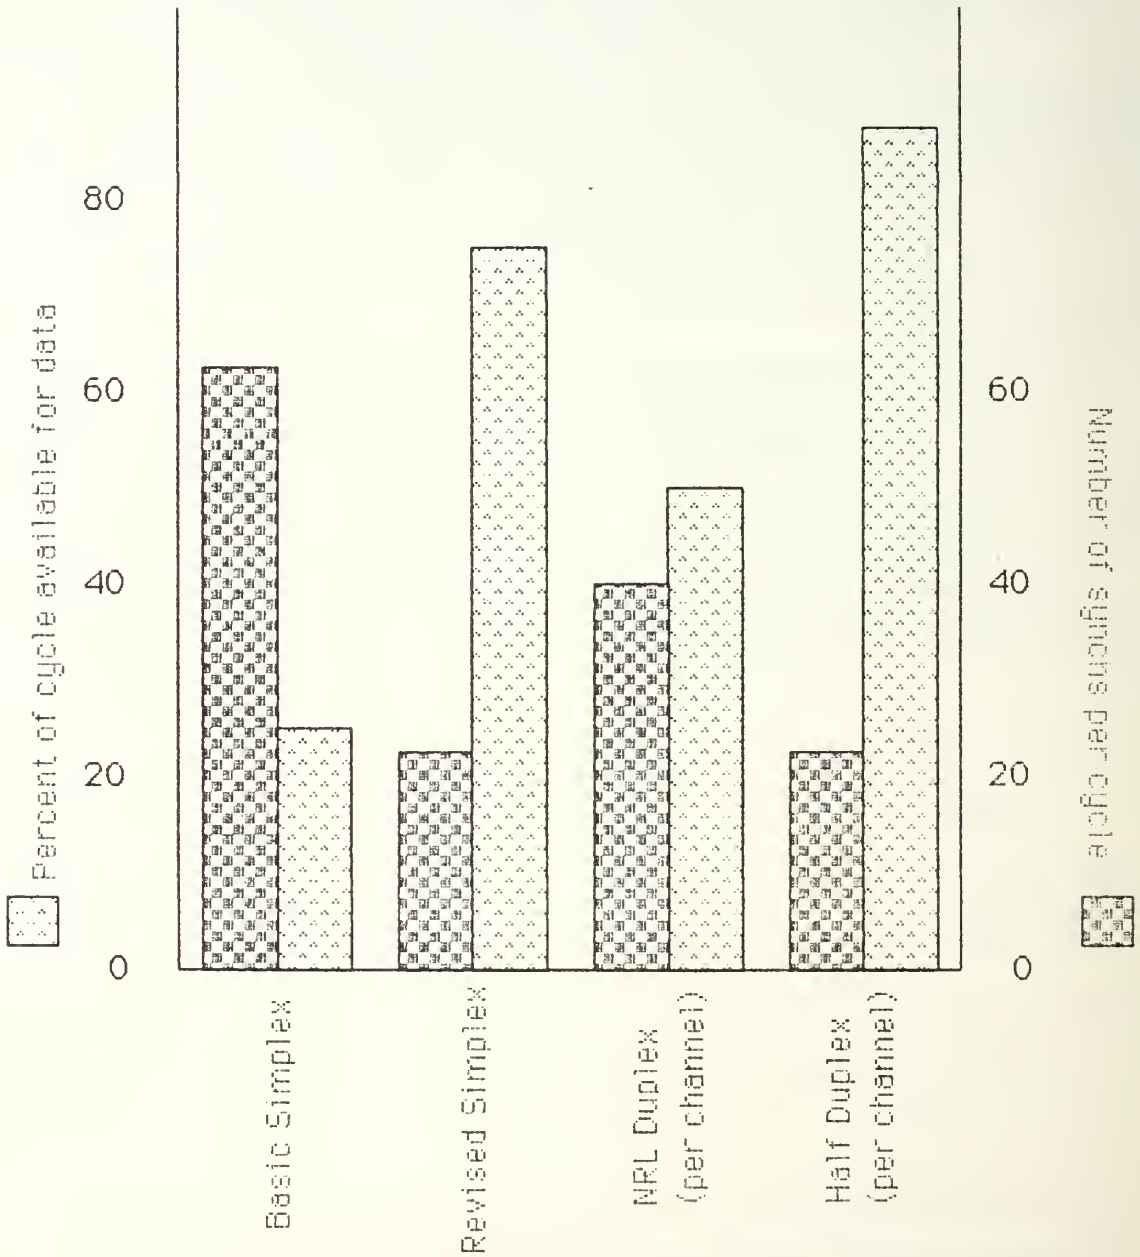

APPENDIX TO CHAPTER--SENSITIVITY ANALYSIS VARIABLES.

For the purposes of our polling circle sensitivity analyses, several variables were stabilized for comparative purposes:

1) cycle time--2 minutes
2) network size--20 ships
3) packet size--256 bytes
4) raw data rate--2400 baud
5) data compression--none (information rate = data rate)
6) forward coding rate--0 (no forward error coding)
7) no ARQ required--all packets received correctly
8) time allowed for silent periods--0
9) synchronization time--1.5 seconds

SENSITIVITY TO VARIABLES.

The values chosen above are arbitrary, but there are reasons why each are realistic. Virtually all of these variables should be operator adjustable and should not be locked into the system at design time.

A two minute polling cycle was used. Why two minutes? There is no empirical reason, but two minutes is a 'good number' for several circumstantial reasons:

1) CUDIXS uses a two minute cycle. There seems no compelling reason to pick another arbitrary number.
2) [Hauser, 84] describe a two minute cycle in their report. The reasoning is incorrect, but it gives a good place to start.
3) Two minutes is a reasonable reaction time. If a ship sends a database query ashore in one cycle, it is not inconceivable that the answer to that query could be returned in the next.

If the communications station operator shortens the cycle time, he will improve access time, and thus grade of service for each ship and and decrease total throughput.

A 20 ship network suits Navy purposes. A number between 10 and 30 means that one network can serve a convoy, battle group, or amphibious ready group. This will offer several advantages in network management. When one ship needs to shift frequencies, all will probably need to shift. Also, if collective addressing in provided in our design, efficiency will improve if all ships in a unit are collected together.

Naturally, throughput and response time will improve if fewer ships are in the network. A number larger than one is vital to our analysis, as the network with only two nodes has no access problems, thus completely masking the important efficiency points.

2400 baud is a reasonable number for current high performance HF modems. There are several manufacturers attempting to sell modems with this kind of advertised performance. If a different baud rate is used, the change will affect all schemes pretty much alike.

A synchronization time of 1.5 seconds was chosen because both the KG-84 and the USQ-83 modem are reported to require 0.8 seconds each. This variable, along with number of ships in the net, are the two critical variables--structure of the network access protocol must properly account for these or efficiency will suffer tremendously--and needlessly. The actual value is less important than order of magnitude--seconds vice microseconds.

The remaining variables are simply arbitrarily fixed in order to perform meaningful estimations. Were the actual values different, each system would be affected proportionately.

Logical Link Layer
Structure and Handling

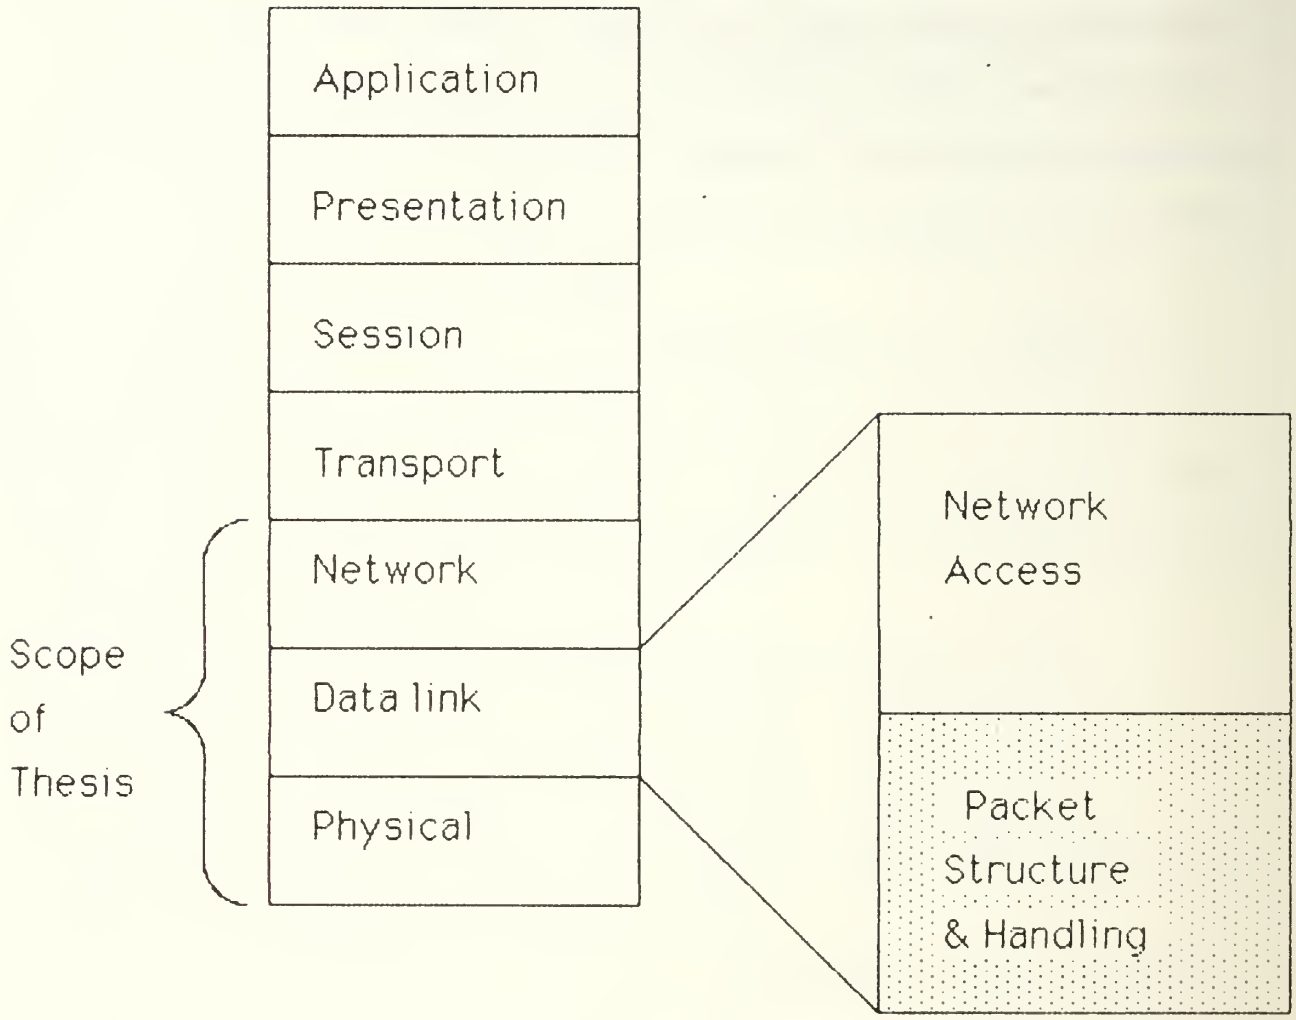

V. LOGICAL LINK LAYER STRUCTURE AND HANDLING

A. INTRODUCTION

This chapter deals with the logical link issues of a sea service communications system. It may be conceptualized as the lower half of the ISO logical link (network access being the upper half).

1. Perspective

This layer in the ISO reference model could not be properly dealt with until the networking and network access problems were dispatched. The feasibility of the architecture in this chapter depends upon the support of the Network Protocol and a network access scheme of the previous chapters.

Unlike the portion of the thesis before this, there are several proposals to perform the functions within the HF band being developed. This chapter is intended to provide overarching architecture rather than a detailed building blocks approach.

2. Other actors

There are other factors which our logical link architecture will accept and deal with.

a. The xisting hysical ayer

Obviously, we cannot greatly change the ionosphere or the way it refracts HF radio waves. So we will live within and take advantage of what we have.

b. The Existing Physical Layer Equipment

Because this thesis is concerned with architecture, the intent is to provide a framework able to accomodate a variety of physical layer equipment such as transmitters and receivers.

1. Existing radio frequency (RF) equipment, as well as planned replacements, should be amenable to the architecture here presented. Once the basic architecture is in place, individual links can be upgraded as 1) the mission requirements demand, 2) as electronic engineering advances allow and 3) as funding is available. These links can be improved without adversely impacting other portions of the system. In particular, the network and higher layers need not be changed. Because the integration of various links is performed at this level, changes to a link are isolated to that link.

This is the basic intent of the layered model approach. Planned evolutionary development of the communications system is envisioned.

2. Because any alternative would be cost prohibitive, link encryption with a synchronizing device such as the KG-84 cryptographic device will be assumed. This factor is less important in this chapter than the last, as the greatest impact of a synchronous physical machine is on network access, not raw data transmission.

3. The Specific Problems

There are two primary problems to be considered in this chapter:

1. HF channels come in narrow bandwidths. Most conventional HF communications takes place over narrowband channels 3kHz wide. The capacity of an HF channel is much less than those of conventional network links.

   In a slightly broader view, there is 27mHz of bandwidth available in the entire HF spectrum. Of that, roughly a fifth is usable between the highest and lowest usable frequencies (MUF and LUF) at any one time (highly variable due to distance and other factors). Of this remaining 4-5mHz, not all is available to the sea services--there are also other users.

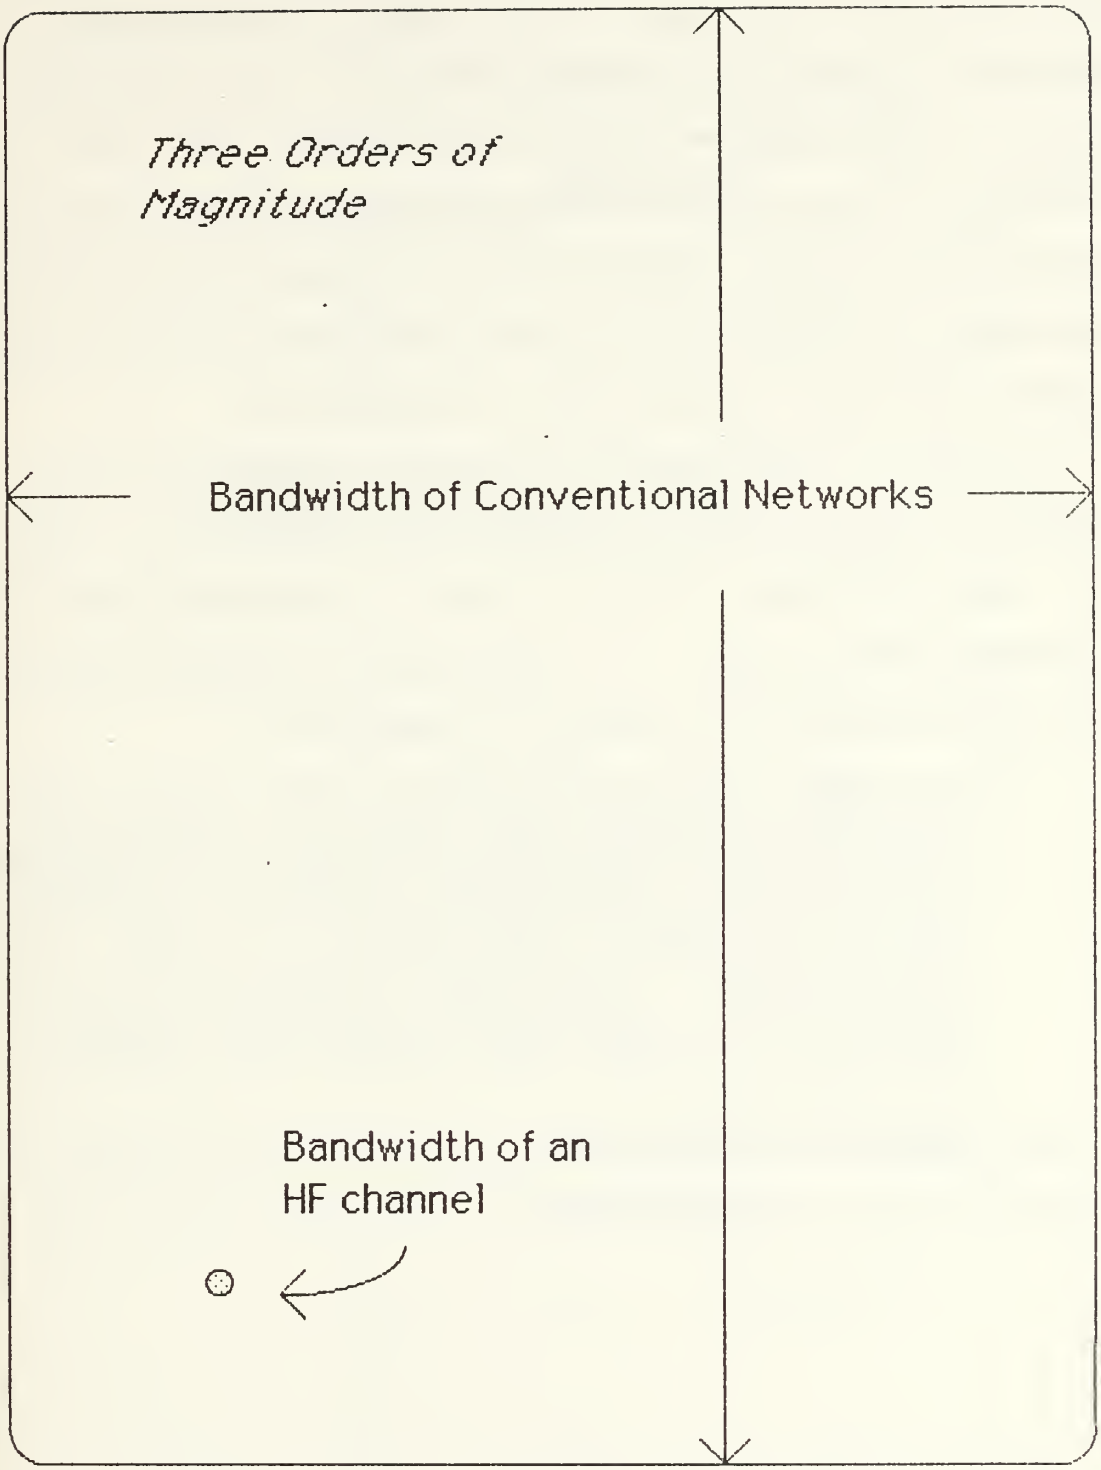
Three Orders of Magnitude
Bandwidth of Conventional Networks
Bandwidth of an HF channel

The illustration should give the reader an idea of the magnitude of the problem. The 'big pipe' cross section represents a conventional network channel with a bandwidth measured on the order of 10^6 bits per second. The 'small pipe' represents the capacity of an HF narrowband channel measured in terms of 10^3 bits per second (after we replace the existing 75 baud modems with state of the art ones that run in the 1k baud-10k baud range and institute data compression).

In short, we can improve our modem equipment and gain roughly one or two orders of magnitude in performance. But even then, that performance is three orders of magnitude less than that experienced ashore. For that reason, squeezing every last bit of efficiency from a channel is important.

2. The other side of the capacity coin is error rate. HF channels are cursed with high noise levels--about three orders of magnitude worse than in conventional network links. Conventional network links ashore experience errors in the 1 bit in 10^6 range (fiber optic is turning in performance in the 10^9 neighborhood). By contrast, HF links have error rates of 10^2 to 10^3. These rates are highly variable depending on a great many factors from sunspots to refrigerators.

The problem is again about three orders of magnitude, but this time in the opposite direction. See the illustration.

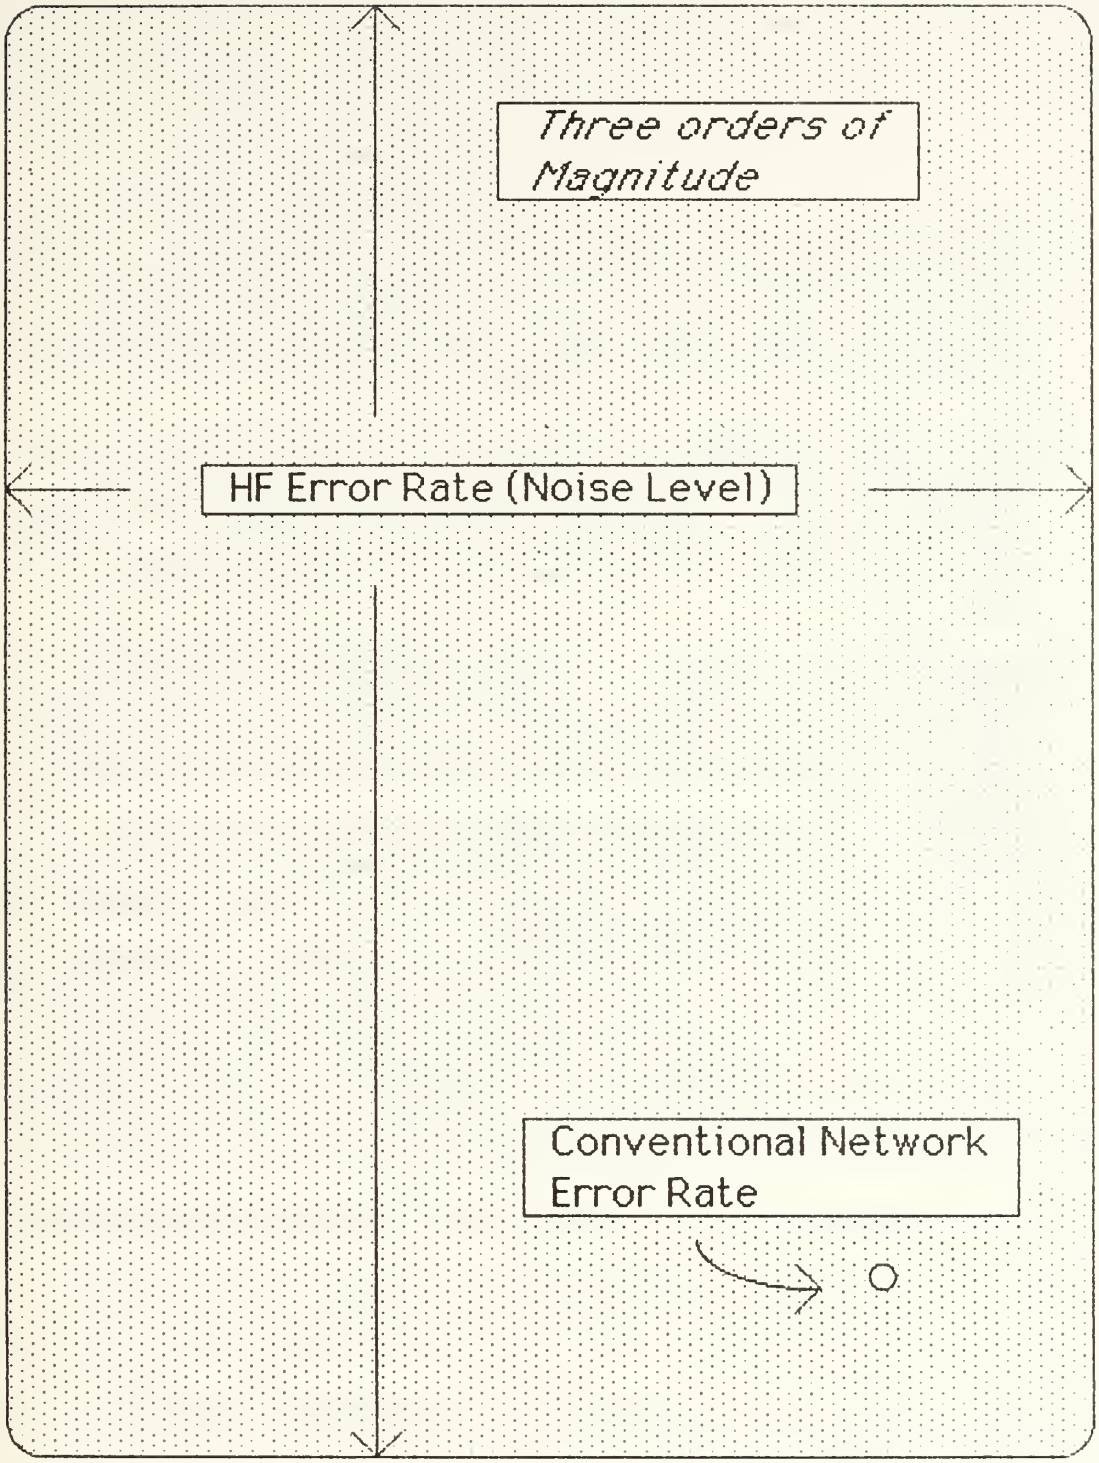
Three orders of Magnitude
HF Error Rate (Noise Level)
Conventional Network Error Rate

Since this thesis is not about HF radio noise, but rather a communications architecture, a detailed discussion will not be presented. But very briefly, the errors will be classified (somewhat arbitrarily) into three groups according to their impact on the communications at the digital level:

1) Burst noise. This is the kind of noise that is caused by lightning. It results in errors ranging from 1-5 bits in length (depending on burst intensity and duration, and on baud rate). In addition to lightning, anomalies in the ionosphere may cause short term fading that results in essentially the same effect.

2) Short Term Fading. Errors that persist longer than a few bits result in the corruption of several bytes of information--or 'holes' in packets. Again, the specific causes are many, but the effect is the same. For our purposes, these are the errors that corrupt parts of packets.

3) Longer Term Fading. The most prevalent cause is diurnal fading caused by the rotation of the earth. The two ends of the link become relocated relative to the ionosphere and a given frequency no longer works. It is necessary to shift frequencies.

   Another cause of fading is severe ionospheric disturbances or SIDs caused by sunspot activity. The only remedy for SIDs is to wait them out and retransmit the data lost during the disturbance.

   The result of this class of errors is packets that are completely lost. Regaining missing packets is the function of the transport layer above the Internet Protocol previously discussed. Consequently, that portion of the discussion is outside this chapter. But the tools to manage frequency shifting do fall within this chapter and will be discussed after the basic architecture is outlined.

The logical link structure presented below attempts to counter each of these classes of errors.

B. ERROR CONTROL IN A BANDWIDTH CONSTRAINED CHANNEL

Automatic Repeat Request algorithms are effective in assuring error free reception of data because they require retransmission of flawed packets. If the capacity of our ship-shore physical layer was adequate to the communications task,

this would be sufficient. But other methods of error control are available which will allow us to make our system more efficient, without sacrificing any effectiveness.

1. Automatic Repeat Request

Our basic ARQ system can be optimized. Messages were broken up into packets in order to optimize packet size to the medium. Since that medium changes constantly, the packet size may need to be adjusted periodically to account for changing error rates. While no balances between errors and packet sizes will be presented, the architecture to control this will be.

2. Forward Error Correction

Forward error correction (FEC) can combat burst errors. Theoretically, most any error could be managed by forward error correction, but its addition rapidly becomes very expensive in terms of capacity sacrificed in order to get it. Consequently, we can practically expect forward error correction to combat burst errors affecting 1-5 bits at most.

Forward error correction is a method of detecting and correcting error without requiring retransmission of the errored data. This link level function has an analog at the physical level: increasing the signal to noise ratio (S/N).

The cost of forward error correction also has its analog in the physical layer. At the data layer, forward error correction involves adding bits to the error stream that allow a receiver to deduce which bits are in error and to correct them. Adding bits to the data stream adds overhead--if one error correction bit

is required for every data bit, the capacity of the channel is halved. The physical analog is narrowing of the bandwidth.

Several texts explain the detailed algorithms involved in adding error correction coding to a data stream. These algorithms are adequate and will not be repeated here [Roden, 82; Lin, 83; Wakerly, 78; Berlekamp, 74].

In order that error correction schemes work properly, they must first detect errors. And indeed they do. But the Cyclic Redundancy Check (CRC) algorithm used in ARQ systems is more effective at pure detection. Its drawback is that it can declare a packet to be with or without error, but cannot determine the location of the error. (Nor can the CRC algorithm tell how many errors are present in a packet--the output decision is binary -- either there are errors, or there aren't.) Forward error schemes are less effective at detection, but localize the error and correct it.

Because bandwidth is precious in the HF environment, forward error correction cannot be simply embraced and applied--throughput suffers too much. Instead the application must be adapted to the error rate (noise level).

The result will be that forward error correction schemes will be likely to detect and correct burst errors that corrupt one or two bits, and will be less apt to handle longer errors. These errors must be left to the error detection ARQ scheme and its associated retransmissions.

A second technique is used in forward error correction schemes to aid in correcting burst errors. Interleaving is used to separate the physical adjacency (the order in which they are

actually sent) and the logical adjacency (where they fit in the message) of bits in the waveform. Thus a burst error that corrupts four bits, for example, can be decomposed into four one-bit correction problems which can be handled individually.

The disadvantage to interleaving is that it increases the overhead expense of reversing the line. Interleaving has the same effect of synchronization preambles as discussed in the previous chapter. This does not rule out interleaving, but it must be implemented with care. For instance, in a half duplex arrangement, a deep interleaving algorithm will have little effect on the shore-ship link, but a somewhat more substantial cost in ship-shore links with multiple ships.

The feedback system associated with ARQ systems can be utilized to adjust the forward error correction rate as well. All that needs to be done is that the receiver package its reception diagnostic information into a control packet which is included with ACKs and NAKs and sent to the sender.

Forward error correction is a useful tool in correcting the short burst errors common in the HF environment. But since it is less effective at combatting fading errors, another tool is needed before we are ready to put together a complete system. This is the majority voting scheme presented next.

3. Majority Voting

Majority voting as a method of error correction is a hybrid of forward and backward error correction. It has particular applicability in our HF environment as a method to control short term fading errors. We visited a variant of this earlier when

discussing assembling packets from various links into complete messages.

The basic theory of majority voting is quite simple. If a packet is transmitted three times (either on 3 frequencies, or on the same frequency 3 times), the result might look like:

```
Wjat6hau Grd =ro;gh#?
eha5 has.Go[ w'ou9ht?
Wh4t *as  od/wr@ug\t?
```

If each bit is 'voted' (vertically in this example), errors that occur only once can be corrected. Since each character in this example was transmitted correctly at least twice, the voting results in a majority in each instance:

```
Wjat6hau Grd =ro;gh#?
eha5 has.Go[ w'ou9ht?
Wh4t *as  od/wr@ug\t?
--------------------
What has God wrought?
```

This illustration is character by character majority voting for illustration. A computerized majority vote system would vote each bit rather than each character, but the theory is the same. Data to be majority voted must be sent an odd number of times (to preclude tie votes) and at least three times. The amount of redundancy must be scaled to the level of noise on the channel.

In most practical uses of majority voting to date, this overhead was fixed--data is transmitted multiple times, whether or not the voting was needed.

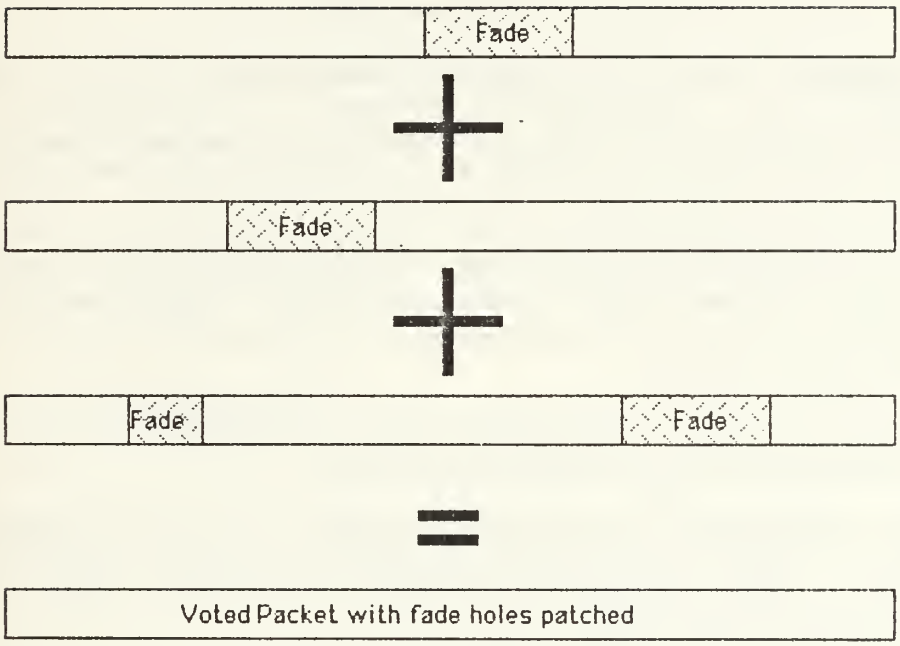

Fade Error Correction with Majority Voting

The following procedure uses a majority voting system within a packet network. Packets are verified within the packet receiver using the standard ARQ technique.

The majority voter only comes into play when the other methods have failed to correct all errors. An example would be where frequent fading on a channel is occurring and the packetizer is able to get partial packets:

a. Each damaged packet, instead of being discarded by the receiver, as in conventional systems, is forwarded to the majority voter. If subsequent retransmissions yield correct packets, the majority voter discards the damaged ones.

b. If three retransmissions occur without getting a correct packet, the majority voter now has the necessary material to work with. It majority votes the three instances and hands a copy of the voted packet back to the packet receiver which rechecks for errors.

c. If the packet is now correct, work progresses to the next packet and the voter dumps its copy.

d. If the voted packet is still damaged, it is retained by the majority voter, as the one closest to being correct and the packetizer gets further retransmissions. The loop is repeated from step one.

e. In response both to the majority voter's inability to synthesize a correct packet and also in response to the bit error rate, the receiver should negotiate 1) an increase in the error coding 2) a decrease in the baud rate, and 3) a decrease in packet size with the sender.

When the monitored error rate decreases at a later time, these parameters can be readjusted to recover some throughput.

This algorithm has some unique advantages. First, the majority voting scheme can effectively 'patch up' packets with fade holes in them--exactly the errors that the forward error correcting algorithm is not well equipped to deal with.

Second, the majority voter does not cause any overhead in the communications system. Packets are only sent to it when they

are received in damaged state and require retransmission anyway. This is an excellent example of computing power being traded against bandwidth.

Third, not all stations in a net need be equipped with a majority voter. It may take more repeats for an unequipped station to receive a correct packet, but a mix of voter-equipped and unequipped stations can work on a net together.

One caution: the illustration is deceptively simple. The voter may be holding a great number of damaged packets in its buffer space at any one time so the buffer may need to be sizable to gain reasonable efficiency. But computer memory is inexpensive and the programming required to implement a majority voter is straightforward.

A tactic that is theoretically similar to majority voting can be used in those instances where the conventional CRC check is, by itself, an inadequate guarantee of message integrity. If the 0.002% probability of an undetected error creeping into a message is too high, the message can be sent twice and compared.

The same majority voting practice, this time called file compare, is used to see if both messages are identical. The probability of a random error striking both transmissions in the same place, causing the same non-detection by the CRC algorithm is practically infinitessimal. Naturally, if the two messages do not compare exactly, a third transmission will provide the necessary voting material to tell which is correct.

This message comparison technique would be implemented at the transport level so a further development will not be done here.

This section has illustrated that conventional ARQ and majority voting can be used in concert to combat the error problems of the HF band. Additionally, forward error correction, in combination, can maximize the throughput on the HF channel with only the bandwidth sacrifices necessary.

We now progress to the subject of data compression, in order to gain as much communication capability as is possible in the channel that is available.

C. DATA COMPRESSION

The object of data compression is to maximize the information content of each bit transmitted across the medium. Several means are available to reduce the entropy of a message. Some different algorithms are discussed in Appendix I, but the tokenization scheme is presented here for use.

1. Tokenization

Tokenization means that the data compressor replaces a character, word or phrase--any bit sequence that has an entry in the token library--with a bit sequence that is shorter. At the receiver, the algorithm is reversed--the token is replaced with the original bit sequence.

2. Why Tokenization

Several advantages drive the selection of tokenization over other schemes.

a. Efficiency

Tokenization schemes have been shown to be quite effective for data sequences that have predictable structures. Source code of computer programs and telegram style military messages are two examples. Efficiencies of 3:1 compression are not uncommon in these instances.

b. Operability Over Several Kinds of Data

In addition to text, graphic images usually have large amounts of white space, data which can be compressed efficiently.

One kind of data that will not compress is a bit stream that has been end-to-end encrypted (E^3). Since the bit stream is random, no significant compression of the data itself will occur. Nonetheless, efficiencies will be gained in compression of the packet structure around the encrypted data.

c. Error Resistance

Since most practical token systems have the static lookup tables (imbedded in the sender and receiver), there can be no loss of the packet containing the lookup table as could occur with dynamic systems.

An unrecovered error occurring in a token will result in an incorrect expansion (possibly no expansion) of that token. But the error will be confined to the phrase that was tokenized. It will not propagate throughout the entire packet. This drawback appears in some other compression schemes.

d. Practicality

Tokenization can easily be performed on a data stream. Some other algorithms are better suited to handling batches of data. For example, Huffman coding generally uses a two pass algorithm that is unsuitable to application to a real time data stream.

D. LINK MACHINES

We turn now to two block diagrams of the logical links. Since we are modeling the links as one way channels, we need senders and receivers. If a communications site has both send and receive capability, then it has two links or the functional equivalent of a duplex link.

First we will discuss the send half of link termination equipment which deals with channel access, packet assembly, and transmission. Following that, a receiver is described.

In the diagrams, several control lines run from the controllers to various components in the data stream. If those components are on the radio side of the cryptographic equipment, a red/black penetration occurs. This is a potentially serious problem with covered circuits. Without the penetrations, the adaptiveness of the links is forfeited. But the penetrations are of very simple control signals, never data. In some cases, the signals are nothing more than on/off settings of a send switch. In others, a simple numerical control is being transmitted. In all cases, these signals can be generated in hardware, so the software verification problems can probably be avoided.

1. A Sender

The data transmission path proceeds from the packet assembler, through a link compression algorithm, link encryption, error coding, modulation and transmission.

a. Packet Assembler

The packet assembler accepts a data stream from the node. The format of this data stream is described in the interface section. The assembler constructs a header and calculates a CRC creating the packet, or envelope that the data is cocooned in for transmission.

(1) Bit transparency. The X.25 standard describes a bit-stuffing algorithm that ensures bit transparency. The packet assembler needs this so that it can transmit any sequence of bits without the receiver confusing certain patterns with the beginning and ending flags. The bit stuffing algorithm is adequate as it stands.

(2) Adapt packet size to error rate. Packet size is controlled by the link control unit and is based on the bit error rate reported by the intended receiver. Only enough bits are accepted from the node to fill a packet. If a push is received before the number of bits authorized to fill a packet are read in, the packet is closed out, and dispatched. The fact that the last packet of a logical message is shorter than the 'standard' is of no consequence since the packet disassembler will be able to recognize both starting and ending flags.

(3) Fragmentation. This operation is very likely to result in packet fragmentation over the HF portion of the net because packet size is probably shorter on this link than any

other in the data chain from writer to reader. For this reason, the Network Protocol contains both message and packet identification material in the header to allow the receiving Network Protocol to reassemble messages.

b. Compression and Encryption

The completed packet then goes through the data compressor and encryption device. For this thesis, the KG-84 is assumed.

c. Forward Error Coding

Forward error coding is then performed to harden the data stream against the expected channel noise level. Again, the amount of error coding added is controlled by the link control unit. This needs to be adaptive because the constricted bandwidth of the HF channel will not support excessive forward error encoding.

d. Modulation

Several manufacturers are building HF modems that handle between 1k and 10k baud. 2400 baud is a common number. Again, like the error coder, the modulator baud rate may be changed to adapt to medium noise.

e. Radios and Antennae

The transmitter and antenna may be conventional HF equipments.

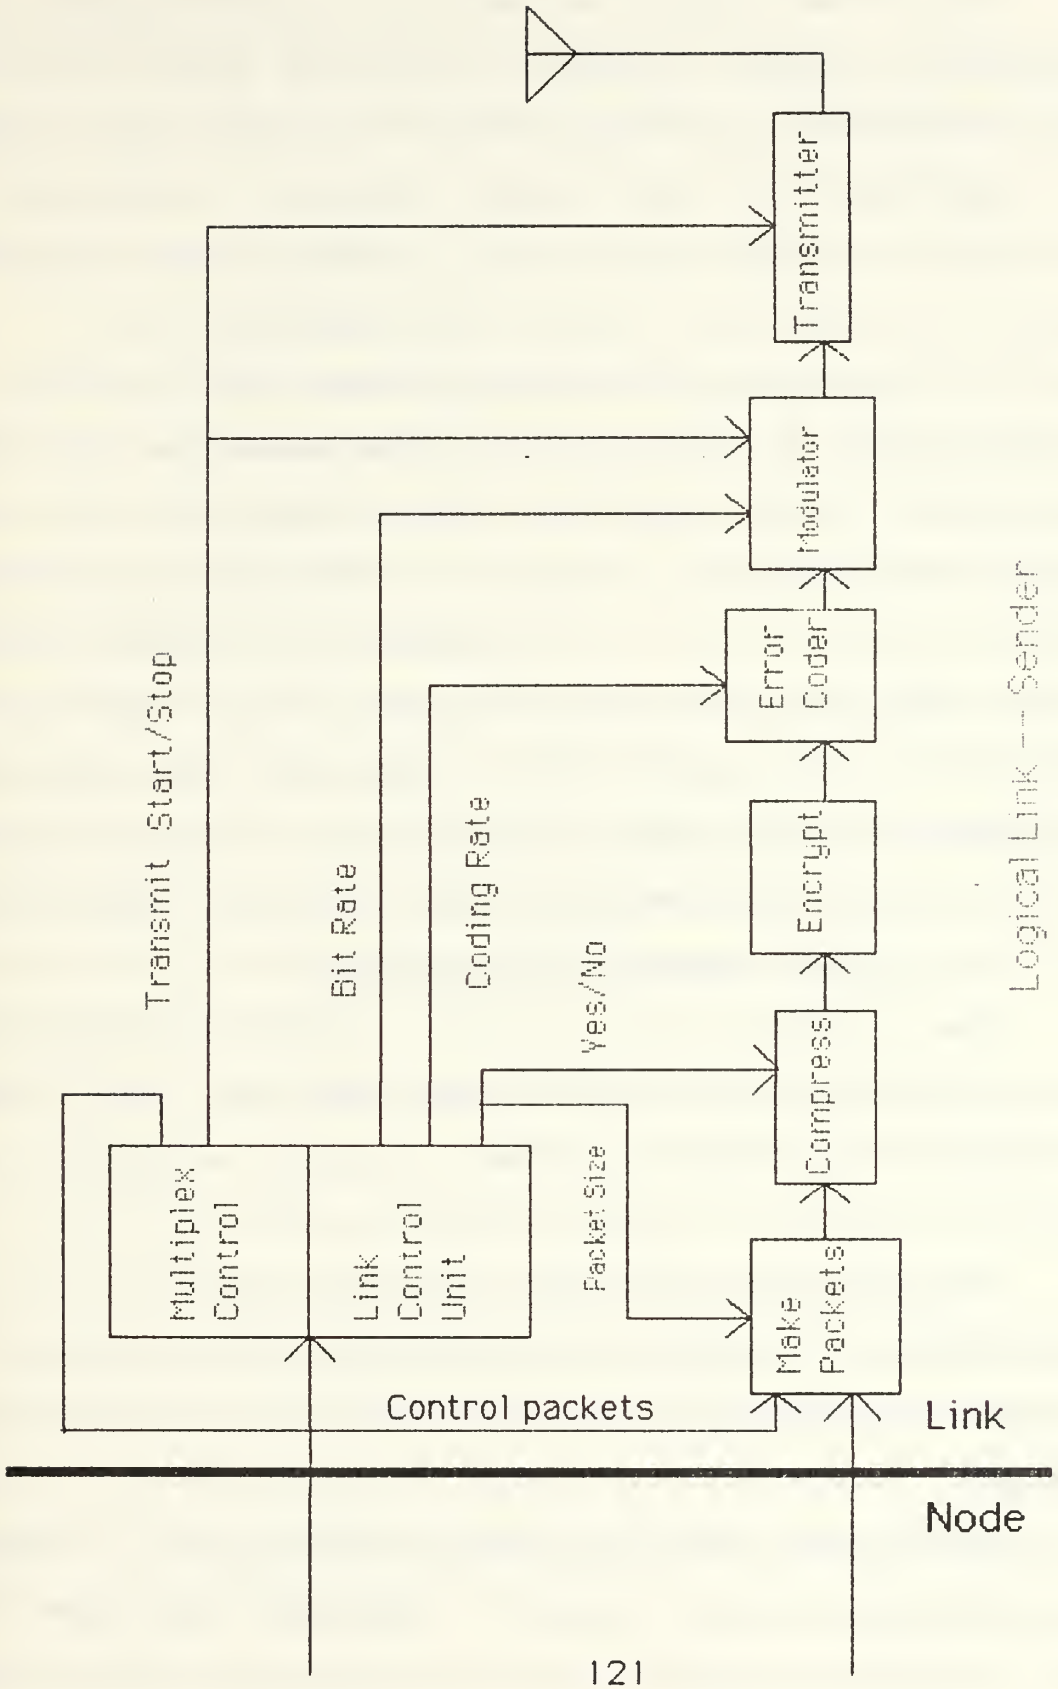
Logical Link -- Sender
Transmit Start/Stop
Bit Rate
Coding Rate
Yes/No
Packet Size
Multiplex Control
Link Control Unit
Make Packets
Compress
Encrypt
Error Coder
Modulator
Transmitter
Control packets
Link
Node

2. Data Stream Flow Control

Data stream control can be handled fairly easily in this pipeline of devices. Assuming a modulator speed of 2400 baud, and a compression efficiency of 3:1, the node might be required to supply data at a rate of about 7k baud. This data must be available on demand, and the rate can change from moment to moment as packet size, coding rate, modulation rate, synchronization preambles and compression efficiency vary. Due to the modest (by computer-to-computer standards) data rate, it is practical to use an RS-232 or RS-422 physical link to hook these devices together. Some sort of flow control such as the CTS/RTS signals in the RS-232 standard will control the data stream adequately (this control signal does not appear on the diagram). It is important that a continuous data stream be presented to the encryption device and all devices downstream, as they will lose synchronization with physically bursty traffic.

3. Packets Awaiting Acknowledgement

When a packet is transmitted, a copy is returned to the network level (node) to await acknowledgement. If the packet is addressed to multiple receivers, the packet is kept in the node's wait list until all receipt packets are in.

If a NAK is received or T1 times out, then the node returns the packet to the assembler for retransmission (the packet competes for precedence with the rest of the incoming data stream).

If multiple links are available from sender to receiver, a retransmission may get routed to a different link port. In this way, a failed or jammed link is automatically routed around.

4. Link Control

The link control unit is not in the data path; it performs the control functions for the sender. This is the most complex part of the logical link layer and will present the greatest difficulty to the developer.

a. Multiplex Control

The multiplex control unit has the function of performing the network access timing described in the previous chapter. In a ship station, the SOL packet will be delivered and evaluated here so the controller can start and stop transmission within the authorized time window.

At the end of each transmission, the controller will determine the queue sizes of pending traffic in the node, compose the status message and append it as the last packet through the packet maker before ceasing transmissions. Consequently, the controller will need to be able to communicate with the node to ascertain traffic queues.

This is also the point where EMCON control should be asserted. EMCON is imposed by sending a 'cease transmission' order to the controller.

b. Working with Other Local Link Terminations

In a downward multiplexed environment, the controller must have knowledge of the HF receivers at the same node so it can inhibit transmissions when data is being received. The link controllers of each link communicate with each other by passing control packets via the node.

The multiple senders procedure is to gang the senders together at the link level. This is a practical solution as it does not require the node to get involved in link functions other than to pass control packets between senders. The controller of one sender becomes the master and uses uses one SOL or cue to control all the transmitters as one.

c. Communications Station Controller

The Net Control version of the multiplex controller must pull more duty than the ship station version. Otherwise, the link termination equipment is identical at either end.

The net control multiplex controller receives entry requests from ships entering the net (via the node). It also maintains a current SOL and a database of the status of all ships in the net. The status message generated by each ship at the conclusion of that ship's transmissions go here.

The controller also has knowledge of its outgoing queue size, coding and modulation rates. Armed with this data, it composes a new SOL for each cycle and delivers it to the packet maker. These control messages are vital to the efficient and responsive operation of the system and must therefore be accorded a precedence ahead of all data messages. These control messages are analogous to the Internet Control Messages in the ARPANET architecture, except that they remain within the network.

5. Link Receiver

The mating half of a send link termination equipment is the receive link termination equipment.

The data path is essentially a mirror image of the sender. The antenna and receiver provide an analog baseband signal to the

demodulator which sends a bit stream to the error decoder. The decoder sends the forward error corrected bit stream to the KG-84 for decryption and signals the bit error rate to the link control unit.

Note that a bit error rate figure can be generated even though the ship may not be receiving traffic addressed to it. In fact, it needn't be logged into the particular net to generate this information--the ship can be totally passive and still gain the data. This thread will be taken up in the discussion of frequency selection where the data will become useful in channel selection.

The error decoder should be able to detect the coding rate from the data stream itself without requiring rate adjustment from the control unit. If an external input is required, the coding rate must be transmitted in a control packet, such as an SOL.

This will make rapidly adaptive rates rather difficult to implement and will reduce the efficiency of the channel even further. The decoder must be able to deal with a data stream that has a variety of coding rates--while this particular link may not require much error coding, some packets such as multiply addressed messages and especially the SOL, must be geared to the lowest common denominator because they are addressed to everybody.

If self-adapting capability is not included, efficiency of the channel will tend to be dictated by the least efficient member of the net.

After decryption, the bit stream is decompressed into its original size and sent to the packet checker (the disassembler half of a PAD). A copy of the unencrypted packet is also sent to the majority voter.

The packet checker reads the address and discards the packet if it is not addressed to this unit. If the address is valid, the packet checker will recalculate the CRC and check it with that made by the packet assembler on transmission. If the CRCs match, the packet is forwarded to the node. If necessary, the packet checker will add the necessary information to the header that was not included in the physical transmission (such as frequency used).

The checker will also generate an ACK packet if full_ARQ is indicated and forward it to the node for transmission back to the sender.

The checker can provide data to the receiver link control unit about the frequency of errors at the packet level. Based on this information, and the bit error rate, the link control unit generates control packets that are routed to the sender negotiating changes to the packet size, forward error coding rate, and modulation rate.

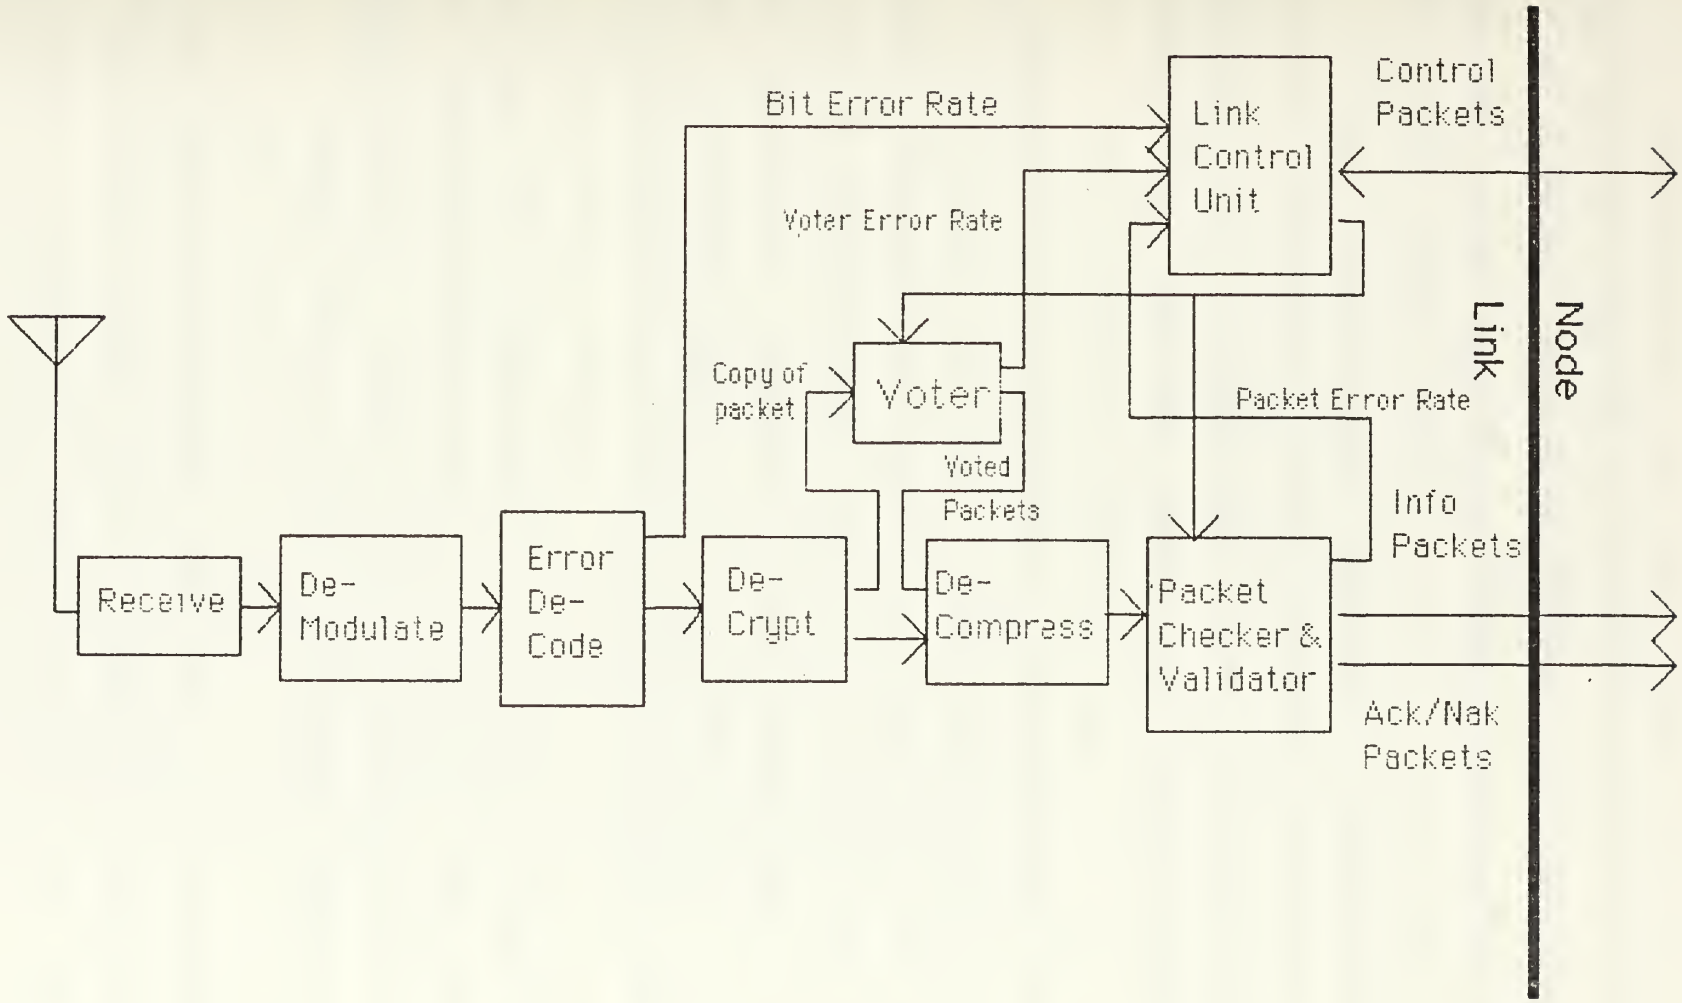

Logical Link -- Receiver

When the packet checker receives an errored packet, it generates a NAK which is forwarded to the node. But the packet is not discarded as in conventional systems. Rather it is forwarded to the majority voter for storage. If the link passes the same packet three times unsuccessfully, there will be three copies of the packet in the voter. The packet can then be majority voted to create a composite packet that will have the best of the three parts. This voted packet is then passed back to the packet checker. If it passes this time, it is treated like any other valid packet and passed to the node.

If the voted packet fails, the expected NAK is generated and the voted packet is passed back to the voter--after all, it will have the least errors of the instances held by the receiver.

This algorithm can cycle until either a correct packet is received, or the system gives up in disgust (a repetition counter times out). This activity will also trigger packet error rate data to the link control unit which will cause packet size to be adjusted.

The majority voter must be located ahead of the decompressor in the data path. Otherwise, the decompression algorithm would compound errors and make different instances of errored packets different lengths. This would make the voting process impossible.

The majority voter has been included in the design partly because we need all the help we can get against the HF noise level. But there are some attractive side effects. Programming of the voter should be fairly straightforward and the hardware involved is quite inexpensive in these post-Silicon days. All

the investment in this form of error control (which isn't much) costs us nothing in bandwidth. Rather, it has the distinct possibility of reclaiming some bandwidth that would be otherwise lost to packet retransmission.

Another interesting observation is that the majority voter presents no configuration control problem. Voters can be furnished to distant ships (such as icebreakers) and omitted for less demanding environments (patrol boats and buoy tenders) without regard for network uniformity.

The control line marked 'dump' is one that will be necessary in our system, but will not be found in conventional systems. Suppose a high precedence message is spread over several packets and most of the message, but not all, is received the first time. It may be some time (seconds at a minimum, hours at a maximum) before the retransmissions patch the holes. The node will have knowledge of the high precedence message in transmission (from the ACK/NAK packets generated). The dump control allows the operator (or node control program) to forward the errored packets to the node so the message can be made as complete as possible, pending corrections. This allows the message fragment to be forwarded to the reader quickly in situations where urgency is initially more important than accuracy. There are few things that frustrate a radioman more than having a high precedence message hung up somewhere and not being able to do anything about it.

E. CONCLUSION

This concludes the discussion of the logical link portion of the communications architecture. This layer of the system must have a twofold objective within the ISO reference model.

First, it must be intimately aware of, and adapt to the physical layer exigencies. In this respect, parts of the architecture are tailored to the bandwidth and noise characteristics of the HF channel. Different channels, such as VLF or satellite would require different logical link constructions.

Second, those characteristics should be insulated from the higher layers in the reference model. The transport and network layers should be general implementations that need not know anything about the physical layers employed. This is the reason for a careful boundary between links and nodes.

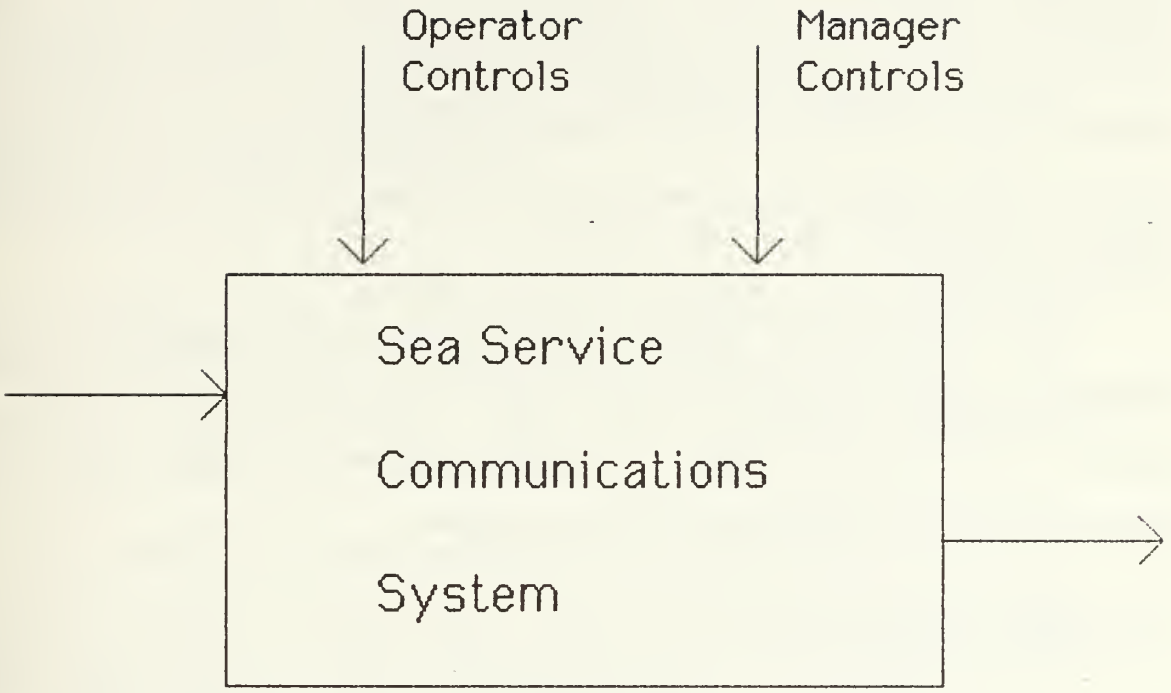
Operator
Controls
Manager
Controls
Sea Service
Communications
System

VI. THE COMMUNICATOR'S VIEW

Thus far, this thesis has dwelt on conceptual models, protocols, SADT charts, ISO references and software engineering problems. But too little attention has been given to the human side of the engineering problem. In this concluding chapter, we look at some of the problems faced by operational communicators and how the communications system outlined in this thesis impacts those people.

By proper documentation and thought to the operator interface, the communications system can be the tool of the communicator rather than the reverse situation where the radioman is an auxiliary to the system.

A. OPERATOR VIEW

Our re-engineered communications system will change the way a radioman does business. Just as microelectronics has ushered in the potential paperless office, this system may bring in the paperless radio central. It most certainly means the paper tapeless radio central.

The most significant shift is that the radioman is no longer in the direct business of sending and receiving messages. Rather his job becomes one of observing flow. This is analogous to the computer room operator who no longer takes your card deck and runs your job for you--rather he keeps the machine running so you can submit your own job.

So what information does the radioman need from a communications system? And how should he be able to influence its operation?

1. Frequency Selection

Long range radio operations in the HF band has always been a challenging business due to the dynamic nature of the ionosphere. Let's look at the tools available within the system to aid a radioman in frequency selection.

First we must envision an operationally sized communications system centered on a communications station. Each node may support several links including HF, VLF and satellite. In particular, in the HF band, there will be several channels in operation.

The communications station will have links set up on several different frequencies. Each link may have any number of ships logged into it and communicating.

The shipboard operator may passively tune up each applicable frequency in succession and 'listens'--gains a bit error rate reading for each channel. He can also get an idea of the traffic loading on each network by examining the respective Sequence Order Listings. The operator will not be reading any traffic (unless the system, perchance, outputs a message collectively addressed to his ship), but he can get activity indications.

The bit error rate indication is closely analogous to reading a test tape. Quality of the channel can be effectively, if crudely, gauged. In many cases, this procedure will be much more effective than use of frequency sounders or other predictive devices as most are based primarily on signal to noise

estimations. The bit error rate 'test tape' is a composite that would also implicitly gauge the other factors such as phase distortion and multipath as well.

Once the frequency has been selected, the operator can simply tune it up, issue an entry request which is sent in a silent period, and enter the net. Traffic is queued at each node, so the system is essentially automatic after that.

If communications are lost--due to a fade, or imposition of EMCON, for instance--the operator may simply repeat the process of finding a frequency. In the meantime, if the ship did not log out, the communications station will continue to send packets 'in the blind' and accumulating a collection of unacknowledged packets. When the ship is again able to communicate, it is relogged into a system, and communications resume where they left off. During the hiatus, both nodes will accumulate unsent or unacknowledged packets in queues.

Naturally, other frequency selection aids can be used if available. It is entirely feasible to include sounder data as control packets that get passed around the network similar to SOLs, ACK/NAK packets and the rest. Use of these aids will improve efficiency and reliability, but they needn't be present for the system to operate effectively.

It is conceivable that the HF frequency selection problem could be automated. This automation would be done by passing control packets that automatically cause receiver and transmitter retuning. But this step should be attempted only after the basic system is in place and demonstrated to be reliable.

If the shipboard operator can use two receive links, he can have one in place, receiving traffic while using the second to scan for a new frequency. Once he finds it, he logs into that link as well. There is no harm in establishing more than one link, even if they aren't required. This way the operator can leapfrog ahead of the diurnal frequency shifts--the new frequency can already be in place and functioning before the old one fades.

2. Operator-to-operator Communications--Service Messages

The existing communications systems depend on the operators' abilities to 'get in each others' shoes' for dependability and efficient operation. This is usually done by inclusion of service messages in the data stream--messages meant solely for the operator at the other end of the link. A second method is by use of a control and coordination net--a separate circuit set aside for use by the operators. Service messages may easily be entered into the system by the operator--just like any other message. The operator should be supplied with a terminal--simply an I/O device connected to the node he is managing--into which he can enter and send service traffic.

If the physical link between the ship and the communications station is reliable, an attended radio central is not necessary. In particular, if the ship is in port with a shoreside telephone tie as the link, full communications can be supported without continuous operator attention.

The major remaining inhibition to an unattended radio room is implementing a message distribution system at each end

(e.g. NavComPaRS ashore) to forward messages to the final destination. Replacing the messenger with clipboard is needed, but was left beyond the scope of this thesis.

3. Observing Flow--Diagnostics

A normally operating node supporting a few links can generate an amount of 'how's it going' data that would smother a mortal quickly if managed in raw form only. Sometimes the detailed data is required, but normally only overall indications of performance are needed.

Overview data from the node should include:

- an indication of queue size. How big is the backlog?
- who is logged into various nets?

Overview data from each sending link should indicate how well traffic is getting to receivers (this can easily be obtained by measuring the sent-but-unacknowledged backlog). This should probably be a negative report--which ships are laboring heavily?

Receiver diagnostic data is fairly simple--data hung up in the system awaiting completion of a message should be dumpable by the operator. By monitoring the ACK/NAK queue, a receiver can be shown to be operating normally.

One diagnostic vital to the function of a radioman is the ability to tell if a particular message has been passed. This query report should be able to take the message identifier and find all or part of the message in question in the node queues or in the storage log. If the message is still queued or partly sent, this kind of status should be retrievable. Also a

collectively addressed message should indicate which ships have acknowledged receipt and which haven't.

More detailed indications of operation will be needed to trace faults and bottlenecks. The operator should be able to obtain additional data by progressing down from one of the reports described above tree-fashion.

For instance, if the shipboard operator is observing difficulty receiving traffic, he can look at all the packet parameters--packet size, bit error rate, congestion (his share of space in the SOL), etc. and determine if a frequency change or additional channel is warranted or if equipment is malfunctioning.

The communications station operator should be able to analyze a backlogged queue and determine who the big users are and balance the load by redistributing the users among various networks and adding links where necessary.

One management option fairly easily available now is the ability to divert queued traffic. If the whole communications system is backlogged, the low priority contents of a queue can be rerouted onto mass storage (tape or disc) and mailed to the ship. Routing queue contents past a message screening board is also quite feasible. The board can decide if a message needs to be sent electrically or mailed and requeues it as appropriate.

The role of the radioman will change dramatically. A great deal of the tedium will be borne by the automated system, leaving the human operator free to concentrate on the less automatable problems of frequency selection and troubleshooting.

B. THE COMMUNICATION MANAGER'S VIEW

In addition to changing the way the radioman views his world, a new system affects the way the next level manager, the communications officer, views his.

The communications station communications officer has new opportunities. From the list of frequencies that he is allocated, he must divide them up among the various users. The basic problem of dividing up limited bandwidth among competing users hasn't changed, but the flexibility is somewhat greater.

First of all, there is no longer a requirement to designate certain frequencies as fleet broadcast channels and others as primary ship/shore. Since each packet is marked with its handling instructions, the channels used need only be channels that work. Fleet broadcast (NAK_only) traffic can be freely mixed with primary ship/shore (full_ARQ) and free broadcast packets (No_ack). The problem becomes a quantitative one of balancing the offered load with throughput and grade of service considerations.

1. Service Classes for Naval Communications

The variant of primary ship/shore known as the itinerant circuit is easily integrated. A subscriber ship will simply not be continuously logged into a net. A ship with periodic traffic will operate without communication until it is required. Then the ship selects a channel, logs in and exchanges traffic. Shore-to-ship traffic for itinerant ships is stored at the node until the ship logs into a circuit. Then the traffic is queued to the link sender servicing the ship. At the conclusion of the

session, the ship logs out and the node resumes storing traffic until the next session (or transmitting in the blind and accumulating unacknowledged packets).

A full period termination is a channel dedicated to the use of a single ship. This is nothing more than our primary ship-shore network with only one ship on the circuit. If that ship's demand exceeds the capability of a single channel, additional channels are downward multiplexed into the system.

How the fleet broadcast is integrated into the system has already been fully described.

A communications shift is easily managed within the system -- if a ship moves from one servicing communications station to another, the pending traffic is simply forwarded to the new communications station and the shoreside routing table is updated.

Because of the inherent unreliability of HF communications, a great deal of redundancy has customarily been built into HF fleet broadcasts. Due to the difference in frequency requirements of several ships all needing the same traffic, the HF fleet broadcast is commonly keyed on several frequencies simultaneously (QLH). Additionally, a rerun circuit is also common--messages are rebroadcast one hour after their original transmission.

Both techniques are easily performed within our packet system--with a considerable improvement in flexibility. Keying several transmitters simultaneously remains a practical option. Two options exist. The first is to have one logical link sender key several transmitters, as is currently the case. The second

is to have one node key several link senders which each support one (or more) transmitters. Rebroadcast can be performed on the same channel if desired (and traffic loading permits)--dedication of a channel to reruns is not needed except to relieve congestion. Indeed, if packets are broadcast at least three times (on any number of channels), the means for a ship to majority vote the data is available, thus decreasing the demand for NAK retransmissions.

One of the more common problems requiring the rerun channel is shipboard power failure. In our system, most of the equipment would be expected to be microelectronics using low power. While providing an uninterruptable power supply for a sender / transmitter is not practical, backing up a receiver is.

A variant of the fleet broadcast is appearing in shore-ship communications today. The Naval Tactical Data System (NTDS) has routinely broadcast data for task force use. Due to the emergence of over-the-horizon targeting requirements and availability of sensors outside the task force, a shore-ship version of NTDS is appearing. The primary difference between this and the conventional fleet broadcast is the elimination of the need for receipt of every message. This type of data is usually refreshed periodically, so it makes little sense to attempt to recover data that will be superseded shortly anyway. This class of packet is also easily integrated within our packet architecture.

The caveat attached to no_ACK packets is that they normally require a high grade of service and are very sensitive to congestion.

C. CONCLUSION

A packet switched ship-shore communications system is practical. The ARPANET experiment which is now being implemented operationally as the Defense Data Network can be extended to cover the special needs of the sea services.

But this extension requires a different architecture at the network and logical link levels of the ISO reference model in order to be effective and account for the unique nature of many channels. That architecture has been outlined in this thesis.

APPENDIX A

THE ISO REFERENCE MODEL

The International Standards Organization (ISO) established an Open Systems Interconnection (OSI) model in 1977 as a basis for people in the communications business to speak the same language. The ISO model is not itself a standard, but is a reference on which standards are built. This appendix briefly describes the model. Further detail may be found in several sources including listed references [Stallings, 85].

DoD's seminal ARPANET work predated the ISO model and evolved a parallel model of four layers, rather than the ISO seven. The differences are semantic rather than functional. For those more familiar with the DoD model, both are mapped.

As you can see, this is the basic form on which the thesis is organized. Some of the standards, such as X.25, turn out to be less than faithful mappings of the model into prescriptions for interoperability. But the the ISO structure as a reference model withstands the unique environment of ship-shore communications quite well.

Layered Architecture Mapping for Communications

ISO terminology	Level	DoD/ARPANET
Application	7	Process / application
Presentation	6	
Session	5	Host to host
Transport	4	
Network	3	Internet
Data link	2	Network access
Physical	1	

Scope of Thesis

A. PHYSICAL LAYER

The Physical layer deals with the transmission of a bit stream over a physical medium. This can further be broken down into two subordinate parts:

1) Analog waveforms in the HF environment. This includes the ionospheric medium, radio and antenna equipment.

2) A baseband layer including modems and link cryptographic devices.

Physical layer equipment generally deals in data streams.

B. LOGICAL LINK LAYER

The Data Link or logical link layer deals with frames of data, error and flow control. This layer is also broken into two sublayers:

1) The packet assembly/disassembly portion of the problem. Metaphorically, this can be likened to inserting a chunk of data in an envelope and ensuring that the envelope is properly sealed and addressed. At the receiver, this layer unwraps the data from its envelope.

2) The network access problem is viewed as another part of the logical link layer. If a packet is to be successfully delivered, only one transmitter can be sending at any one time.

The Network layer deals with establishing, terminating and maintaining connections. In this thesis, a great deal of effort is spent here integrating a heterogeneous collection of data link and physical layers into a complete network.

C. NETWORK LAYER

We subdivide the Network layer into two parts as well:

1) The lower half is concerned with the network level acknowledgement system developed in Chapters II and III.

2) The upper half of the Network layer is served by DoD's Internet Protocol and is primarily concerned with integrating different networks into larger communications systems, sometimes referred to as catenets.

DOD lumps all of the lower three layers that deal with a single network into one Network Access Layer. This places the dividing line between the Network Access and Internet layers at the point where a Local Area Network (LAN) connects to other networks.

This particular dividing line is of use to us because that is the lower boundary of the current Internet Protocol. For most conventional networks (e.g. X.25), the lower half of the Network layer is vacuous. This unused portion is made to order for the Network Protocol that we need. This thesis accepts the current Internet Protocol specification and builds under it.

D. TRANSPORT LAYER AND HIGHER

DOD's host-to-host layer provides some of the functionality of both the ISO Session and Transport layers, probably because the accepted protocol and currently operating software is the ARPANET originated Transport Control Protocol (TCP). The TCP documentation actually feathers into the Presentation layer because it specifies things like File Transfer Protocols that are part of that layer, but generally specified as part of the package required to make a TCP at all useful.

The function of the transport layer is to maintain logical continuity and to control errors such as missing packets and messages. This thesis uses TCP (and higher layers) without modification--we are able to successfully construct an architecture under it.

APPENDIX B

STRUCTURED ANALYSIS AND DESIGN TECHNIQUE

Structured Analysis and Design Technique (SADT) is a methodology for systems engineering and complexity management.

In particular we have a means for implementing top down development. A system is developed in tiers or levels. Level 0 is the highest level. The leading illustration to Chapter 1 is a Level 0 illustration of a communications system.

Parenthetically, whatever subjectively constitutes a 'system' is designated as Level 0. Another designer, perhaps of a weapons system, might consider the communications system of this thesis as just one (e.g. Level 2) component of his system.

A system is represented as a box. Inputs are entered at left and outputs exit from the right. Controls exercised on the system are entered at the top.

Each box can be considered a 'black box' at its particular level. At Level 0, the entire communications system is a black box. This is the view of the ship-shore-ship communications system that we wish to present to the non-communicator. He puts messages in and gets messages out without concern about how they are handled within the system.

At the next level down--Level 1--the Level 0 box is expanded to a series of boxes and interconnections. Some of the detail that is hidden at the top level is now shown. Each of the subordinate boxes now exposed may be further expanded as necessary. The Chapter II frontpiece illustrates this expansion. Figures 2.4 and 2.6 are also Level one expansions showing the differing interfaces.

The link machines (sender and receiver) in Chapter V can be considered Level 2 expansions of the 'Process & Transmit' and the 'Validate & Process Received Packets' boxes of figure 2.6.

Our HF design problem is a complex one and it is at once difficult and necessary to both understand the broad sweep and deal with the necessary details.

Since inputs, outputs and controls are specifically represented, this methodology is very useful for conceptualizing and understanding interfaces and system integration.

APPENDIX C

AN INDUSTRY SURVEY

Every development project should contain an industry survey, and this thesis does. As the thesis developed, and particularly as the network network level acknowledgement idea was accorded its importance in network integration, the survey became increasingly irrelevant to the central thrust of the thesis. For that reason, it is presented here as an appendix, rather than at the beginning of the thesis as an obstacle to be passed before getting to the vital parts.

This appendix describes several communications systems and concepts that bear on a packet switched, high frequency communications system suitable for the sea services. Additionally, one system described--the TNC--has direct applicability in the National Communications System in which both the Coast Guard and the amateur radio community are participants.

None of these systems is wholly satisfactory for our sea service needs. But each has some strong points worth investigating.

Approach. The first model described is the existing HF, ship-shore equipment setup and operational procedure. This anchors the discussion to existing reality and provides a conceptual foundation on which to build. While the existing system is considerably older than the layered architecture ISO model, it can be mapped onto the model after a fashion.

The file transfer protocols available to microcomputers is covered next. The basics of packets are employed in XModem and Kermit.

The third description is industry generic. Covered are the standard architectures used for conventional Local Area Networks (LANs). Topologically, a ship-shore HF network is a LAN, transoceanic distances notwithstanding. Therefore a discussion of these systems is useful as it relates the conventional computer-to-computer communications business to our particular topic.

The fourth model also exists today. It is that used by the amateur radio community for packet radio. This model is of limited usefulness in the naval communications architecture, but is a very useful point of departure. The amateur radio solution can be directly used in some situations in the military and disaster recovery situations without change. In others, it constitutes an effective quick fix pending a more comprehensive installation.

A next two models come from the Navy. First is the Common User Digital Information Exchange System (CUDIXS) which is the satellite communications system used for fleet communications. While that system cannot be directly ported over to the HF band, its protocols are useful in our study of network access.

Following is the Link 11 (Naval Tactical Data System) for intra-task-force data communications. Again, the protocols are not directly applicable to our ship-shore network, they are useful as a starting place in managing the network access problem.

The final system surveyed is the Coast Guard's SCAMP project designed to implement ISO physical and data link layers. One intent of this project is to develop a datalink capability that can be used in the HF band.

A. CURRENT HIGH FREQUENCY SYSTEM

The illustration following is a block diagram of one terminal of an HF network. The existing message switched system can be roughly mapped onto the ISO model as the illustration shows. The mapping can only be rough because of the wide gaps occasioned by equipment installations that predate the ISO model by decades.

The radio and antenna form the physical link components of the communications system. The modem forms the interface to digital signals and forwards a data stream to the cryptographic and teletype machines. The teletypewriter with its built in peripherals (tape reader, printer, punch) is both a level 2 component and an application layer component. The human operator fulfills some of the Logical Link and all of the Network level functions as well as all higher ones.

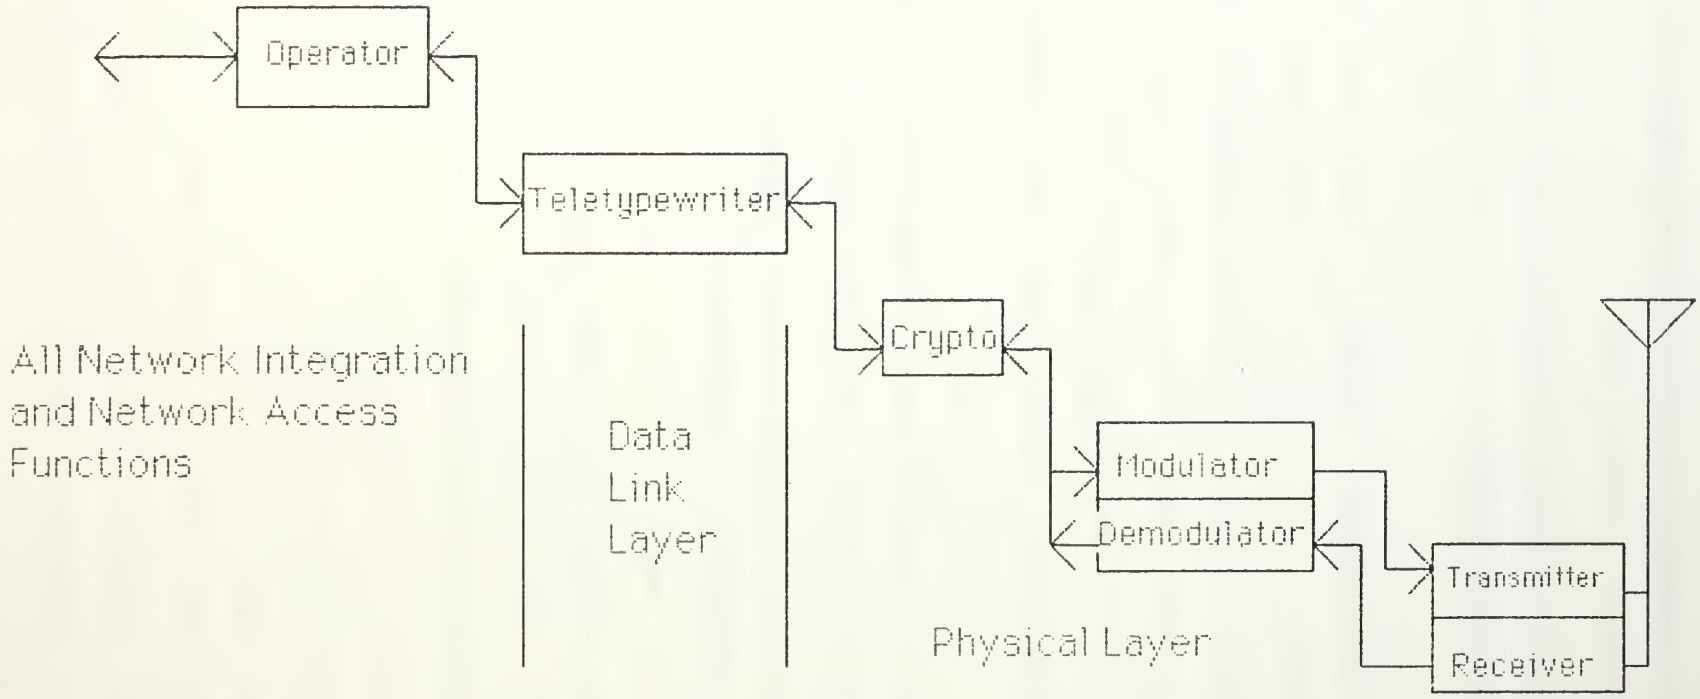

Existing HF System - Shipboard View

Because the long-haul ship-shore communications system itself does no internetworking, it is topologically a star configured local area network, distances involved notwithstanding. Each ship station communicates with the network controller or server--the net control station--the communications station.

In this thesis, the term communications station has consistently been used to denote this network hub. There is no particular reason why the basic architecture could not be used for Intra-Task-Force communications with a unit of the task force serving as the network controller. Indeed that is the organization of the Naval Tactical Data System which is described below.

Another convention of terminology used in this thesis is that all subscribers in the network are termed 'ships'. A node communicating with a communications station could be a ship, aircraft, submarine or served shore station. For shorthand, these have been called collectively 'ships' in the thesis.

Reliability. The existing HF system has no error control inherent to the system save the operator reading incoming messages and servicing for errors. Other than the message 'reading right' there is no assurance of correctness. Reliability of the HF teletypewriter circuits is better than the CW systems they replaced in the 1950s, but not adequate for many of today's needs. Ashore, the NavComPaRS validates each address in the header and queues any incorrect messages for manual correction. But the body of the message itself is not checked.

A second aspect of reliability is that of maintaining the physical link. This is generally considered a maintenance function in other networks, but is an operator function in the HF band. Maintaining a usable frequency between the two units in a link can be a full time job for the radiomen on watch when the distances are great. The complexities of maintaining several links for all the ships in a net require several techniques on the part of the communications station totally alien to normal computer network communications. As we have seen, this task will be somewhat easier with an automated system, but the basic difficulties will remain.

Traditionally, the task involves consulting objective sources such as MUF/LUF tables and frequency allocation tables to get a ballpark frequency. Then the operators exchange test tapes on several different frequencies until a usable one is located. Maintenance of communications involves continued use of test tapes, some rules of thumb (up during daylight, down at night) and a healthy dollop of experience.

Recently, more objective aids such as computer aided forecasters and ionospheric sounders, have become more readily available. These have the potential to offsetting the subjective parts of frequency selection. These tools, at this level, require that the operator integrate the data gained manually into his frequency management task. The job is easier, but only slightly less manual.

The remaining challenge will be to integrate use of these diagnostic and predictive aids into the system.

Most ship-shore HF teletypewriter circuits operate at 75 baud (100 words per minute). As both Navy and Coast Guard missions become steadily more data intensive, this transmission rate is clearly not adequate to the need.

Network access is a manual process. Part of the labor can be offset by use of fleet broadcasts, but the system remains manual. Transmitters take turns, with turns being determined by the communications station radioman who has a manual allocation scheme. This algorithm contains rules such as 1) handle highest precedence traffic first, 2) handle ship-shore before shore-ship traffic, 3) service distant ships before nearby units. The basic scheme is a reservation system.

This system is known in the communications community as a directed net. A free net corresponds to a carrier sense multiple access system similar to that described below in commercial LANs. Free nets are not practical in long haul systems because not all ships can hear each other.

It is this network access algorithm that must be embodied in the link multiplex controller.

The existing system is secured by link encryption. It has some instabilities. All stations in a network have the same keylist and receive all traffic whether it is addressed to them or not. This requires all communications centers to be cleared to the highest level of traffic carried. The existing system also generates classified paper waste at each communications center, thus increasing the opportunities for compromise. In the system described in this thesis, this classified waste problem is managed because the logical link only outputs messages

addressed to that ship. This is not tamper proof, but it decreases the potential for accidents.

Data Transparency. The existing system is designed to handle telegram style message traffic and nothing else. Facsimiles and voice transmissions need separate circuits and separate equipment. Since the new system is bit transparent, the contents of a message makes no difference to the communications system--it transmits a bit sequence without regard to the format of the message contained.

In another dimension, full_ARQ, NAK_only and no_ack traffic must be sent on separate channels. Our system allows a free mixing of this traffic because each packet is individually recognized.

Considering some subsidiary considerations, maintaining the old teletypewriter equipment is expensive compared to current generation equipment. Operators are tied to repetitive tasks that can be automated. The old printing equipment is noisy. And in the baseband level, the existing equipment suite makes it very difficult to update the architecture: nothing can be changed until it all changes.

One of the goals of this thesis was to introduce enough modularity to the system to allow incremental updating--alias pre-planned product improvement. Hopefully, the inflexibility of the system to updates can be avoided.

B. MICROCOMPUTER DATA TRANSMISSION

Microcomputers have given rise to data transmission capabilites. Most systems are designed to transfer data over telephone circuits.

The original modem program for microcomputers was written and placed in the public domain by Ward Christiansen. His XModem protocol provided the basic packet structure with a rudimentary header and a following checksum for error control. XModem fixed packet size at 256 bytes and used a Stop & Wait ARQ protocol.

As modem usage increased, the desire to improve file transfer capabilities also increased. The protocol has evolved (in both XModem and also YModem) to support of 1Kbyte packet sizes and sliding window (Go Back N) protocols. These improvements were occasioned primarily by the advent of 1200 and 2400 baud telephone modems, largely replacing 300 baud modems. 256 byte packets and Stop & Wait required that the phone line be turned around repeatedly and rapidly causing a considerable loss of the speed advantage that was supposed to come with a higher baud rate modem.

Concurrently, the SQueeze and UnSQueeze programs using Huffman coding became common for data compression to cut down on phone costs relative to data carried. See the appendix on data compression for further discussion.

The fundamental improvement of the Christainsen protocol over simply placing the computer into a 'copy' mode and listening on the line was error control.

Using packets as parts of the whole file is a decoupling of logical and physical messages common to all packet systems. This

means that an error causes retransmission of the faulty packet, not the entire message. Use of a checksum for ARQ assured verification of messages received.

Kermit. While microcomputers were readily communicating with the XModem family of programs, micro to mainframe communications turned out to be a more difficult problem. Many mainframes were not amenable to XModem implementations. Further, several have strange limitations such as 7 bit Input/Output channels instead of 8.

This led to development of the Kermit protocol. The basic conceptual notions of Kermit are the same as XModem--packets and ARQ. The difference is in implementation. Kermit uses a more general architecture that can accomodate the quirks of a great many computers (such as prefixing which allows an 8 bit logical byte to be accurately represented in two 7 bit physical bytes).

XModem and Kermit implement logical layer protocols. They use the telephone system both as the physical layer and to perform the network functions of call placement and disconnection. Neither protocol uses any addressing, but rather relies totally on the circuit switched network capability of the telephone system.

These useful programs have lessons for us, but are incomplete for our ship/shore system.

C. COMMERCIAL LOCAL AREA NETWORK SYSTEMS

Commercial LAN systems have two basic parts of use to us. The first is packetization, the second is medium access.

1. Packetization.

A packet is a physical unit of data--several packets glued together in the correct order make up a logical unit of data, a message. Packetization decouples the communication problems of dealing with packets from the information management problems of dealing with messages. Once the problems are decoupled, they can be attacked separately.

In this respect, the packetization in a LAN is similar to that done by the XModem/Kermit programs. The difference is in the fact that a local area network is like a 'party line' so packets must have addresses attached for routing where XModem and Kermit use the circuit switched telephone system to do the routing.

A packet is a fairly simple unit:

- a flag byte. This tells the receiver where the beginning of a packet is in a data stream. X.25, as an example, uses the unique 01111110 byte as the flag byte. Bit stuffing is used to prevent that pattern from showing up anywhere else in the bit stream.

- 'from' and 'to' addresses. This allows a packet to be recognized by the intended recipient and rejected by all other stations on the network. Similarly, acknowledgements to be returned to the sender can be correctly addressed by using the 'from' as a return address.

- some overhead information such as timestamp, data type, etc. Important for us in this control data is the packet serial number which allows the receiver to sort packets out, check for missing pieces, and reassemble complete messages.

- the data itself. The only alteration is the bit stuffing algorithm to prevent the random occurrence of the flag byte pattern in the data. Since the bit stuffing is reversed at the logical link level at destination, it is transparent at all higher layers and the data out appears exactly the same as the data in.

- a checksum of some type. The transmitter calculates this upon packet transmission. The receiver recalculates the checksum and compares it with the transmitter's. If they agree, the packet received is the same as what was sent. X.25 uses a Cyclic Redundancy Checksum which has a very low probability of incorrectly judging a flawed packet as correct.

  This allows the fundamental Automatic Repeat reQuest (ARQ) procedure which is sometimes known as backward error correction (BEC). If the CRC match fails, the packet is NAKed and resent.

- a closing flag byte (which may also be the leading flag byte of the succeeding packet). In X.25, empty space between packets may be filled with successive flag bytes to keep the data stream moving at the physical level so equipment does not lose synchronization.

While the details of packet structure vary among systems, the basic structure remains.

2. Medium Access Procedures

Local Area Networks (LANs) come in four primary flavors of medium access protocols, which is the crucial point in making an HF system operate. These protocols include:

- Collision Detection or Aloha systems.

- Carrier Sense Multiple Access with Collision Detection (CSMA/CD), sometimes known as Listen Before Talk or Listen While Talk.

- Token systems.

- Central reservation systems.

Collision detection is the simplest procedure to implement. Whenever a station has data to send, it sends it. If the packet arrives successfully at its destination, it is acknowledged and the system proceeds to the next packet.

If more than one sender sends at the same time, a collision occurs and no data arrives correctly. As a result, either a NAK or no response at all is generated. These symptoms signify a collision and the senders wait a random amount of time and retransmit. Some efficiency can be gained by slotting the system --a clock synchronizes transmissions so that collisions will only clobber packets in one quantum, not two.

Collision systems suffer two serious drawbacks:

1) They are very wasteful. They reach capacity at a throughput rate that is a fraction of the bandwidth capacity. We cannot afford this waste in the HF radio spectrum.

2) They are unstable. As offered traffic approaches the maximum thoughput rate, the number and rate of collisions increases. Collisions beget more collisions and the system chokes and crashes. In commercial systems, the instability is avoided by building enough bandwidth into the system so that it never occurs. This inability to operate under stress cannot be tolerated in a military system and the bandwidth cannot be sized to avoid the problem.

CSMA/CD. This system is popular because it is fairly easy to implement in hardware. Such systems are quite efficient in applications such as linking several computers as small as microcomputers to disc servers and printers.

CSMA/CD systems are effective where the bandwidth available is large relative to the traffic offered. When computers can be strung together with coaxial cable or optical fiber, this is usually the case. CSMA/CD systems suffer three major problems that make them unsuitable for our use:

1) Since nodes gauge their permission to transmit on whether or not they can detect the carrier from another station, the whole system rests upon the assumption that everybody can hear everybody else. In the HF world, this isn't the case.

2) They are incapable in themselves of coping with a station monopolizing the network. If one terminal continues to transmit continually, all other stations are inhibited. Military networks require the ability to implement an equitability system so that stations may break in with high precedence traffic.

Token. Token systems involve a station gaining the 'token' which constitutes permission to transmit. When it is finished, the token progresses to the next station on the net which then takes its turn.

This method allows a high throughput--efficient use of the capacity available. The primary drawback is the same as in carrier sense systems--each node must be able to hear the station passing the token to it. See the Link 11 discussion below for further discussion.

Central reservation systems are not widespread in commercial local area networks. The algorithms available come mostly from military systems and are described in further detail elsewhere.

Central reservation systems generally require a star topology and are vulnerable to hub failure--if the communications station fails, the entire network crashes. On the other hand, failure of any other single node will not affect the rest of the system as is the case with token systems.

Standards. Packetization has proven useful enough that standards have emerged. Probably the leading standard is the X.25 standard which is described by the International Telegraph and Telephone Consultative Committee (CCITT) and published by the International Telecommunication Union (ITU). X.25 is used as a reference in this thesis because it is the network connection of choice to the Defense Data Network (DDN).

Parallel, and nearly identical standards have been published by the American National Standards Institute (ANSI) and for government use as Federal Information Processing Standards (FIPS). Additionally, some standards are republished--and often changed--as Military Standards (MIL-STD).

The ARPANET Transport Control and Internet Protocols have been formalized as MIL-STDs (1778 and 1777 respectively).

D. AMATEUR PACKET RADIO--THE TERMINAL NODE CONTROLLER

Amateur packet radio proved to be a very interesting area with many lessons applicable to the ship-shore problem of this thesis. An initial hypothesis was to adapt and expand the amateur work to the need. Ultimately, this was not practical, but it was a very useful place to start the thesis.

The work has been pioneered by the Tuscon Amateur Packet Radio organization which designed and produced the core hardware and software for hams to incorporate with their rigs.

The central equipment is the Terminal Node Controller (TNC). The TNC provides the data link level functionality. A continuous data stream from the microcomputer (terminal) is bundled into packets by the TNC and transmitted. The TNC has an integral modem so it physically fits between a terminal and a transceiver.

The packet protocol used is AX.25 which is a variant of X.25. Since X.25 lacked several features at the data link level that were necessary for amateur radio, AX.25 added them--a stretching of the existing protocol. Chief among these features is balancing of nodes. X.25 designates a central node as a master station (elsewhere in this thesis designated net control

station). This is the Data Circuit-Terminating Equipment (DCE). All other participants in the net are slaves--Data Terminal Equipments (DTE). AX.25 makes no such distinctions among nodes --they all have equal status and are referred to as DXEs--Data Switching Equipment. For some military applications, this balanced protocol makes sense. For our HF system it isn't necessary.

The TNC is a Level 2 machine and is designed as a precursor to a Level 3 system that will include it. Some higher level considerations have been built into TNCs, but the implementation is incomplete as of this writing. Level 3 procedures beyond the functions grafted onto the TNC level 2 function rely on operator discipline.

The specific hardware is very straightforward. A TNC is essentially a general purpose microcomputer with a specialized input/output device (the modem and radio controller) added. The specialization comes primarily from the software.

TNC's use the ARQ technique of X.25, so the basic error correction is incorporated.

Performance. The modem integral to the TNC operates at 300 baud, frequency shift keyed (FSK), asynchronous transmission. While this is an improvement over the existing 75 baud of ship/shore circuits, it is less than that feasible. (The FCC currently restricts US amateurs from exceeding this speed in the HF band.) Hams outside the United States are reportedly experimenting with transmissions in the 600-1200 baud range with

good success. This performance is gained within the same 5kHz narrowband channel used for military HF communications.

Simplex operation. TNCs are designed for amateur radio operations which normally run on a single frequency. X.25 protocols require that packet acknowledgements travel along the same path in the reverse direction. This means that the channel must be 'turned around' frequently. Since data transmission is asynchronous, this is tolerable.

In more recent updates to the TNC, multiple connections are possible. TNCs rely on CSMA/CD for network access. Even in the 2 meter (VHF) band where carrier sense is more effective and TNCs operate at 1200 baud, multiple logical connections on one physical link results in a significant decrease in throughput.

A TNC based system would be very useful at low cost in a variety of governmental situations:

- disaster recovery situations. Because of the cost and lack of encryption restrictions, TNC's are readily available to hams who have a history of aiding in reconstitution efforts.
  The National Communications System organization should capitalize on this level of utility. As the amateur radio community itself is very concerned with standardization and interoperability, there are advantages to using their developments rather than inventing new ones if they can be made applicable.
  Unless traffic analysis vulnerability is considered a major problem, the use of TNCs and end-to-end encryption is entirely practical and much lower cost than the more involved construction needed in a ship-shore system.

- fixed networks where link encryption and high speed are not important. One example would be the HF administrative and operational nets joining Loran stations. Unlike our worst case ship-shore model, units do not enter and leave these networks. Reliable operation in the hands of non-radioman operators is useful and TNC's offer a very cost-effective step in this direction.

- commercial ship-shore-ship communications. Some subset of the naval model should be formulated as a merchant fleet, low cost, interoperable model. The existing TNC fulfills the basic requirements.

TNC equipment is readily available and at reasonable prices. At this writing, a TNC can be obtained for less than $200. In many situations, a quick fix may be desirable pending development of the more comprehensive system developed in this thesis. Given the costs involved, the investment can be written off very quickly.

E. NAVY CUDIXS SYSTEM

The Navy satellite communications system is subdivided into several user communities. This section briefly describes one system, the Common User Digital Information Exchange System (CUDIXS). This channel is generally used by the surface navy. Other channels are listed under CUDIXS in the glossary. The architecture is similar, so one description will suffice.

The part of the CUDIXS architecture that is useful to our HF system is the network access scheme. This algorithm is borrowed for use in our upward multiplexing description in the thesis. The net cycle provides a central reservation scheme that keeps ships from jamming each other with simultaneous transmissions.

The system operates by the communications station sending its traffic followed by a Sequence Order List (SOL) indicating when a ship's turn is coming up. Each ship broadcasts in turn.

Blank spaces in each cycle are interleaved in each cycle for new stations entering the net to send their entry request packets. These are called Random Access Time Slots (RATS). RATS

access is contention based--the CSMA/CD model. (While the algorithm is borrowed, the acronym is not--in the thesis, these are silent periods).

Equitability is enforced by allowing each ship a maximum number of blocks in each cycle.

CUDIXS is a central reservation scheme that efficiently uses bandwidth within the channel.

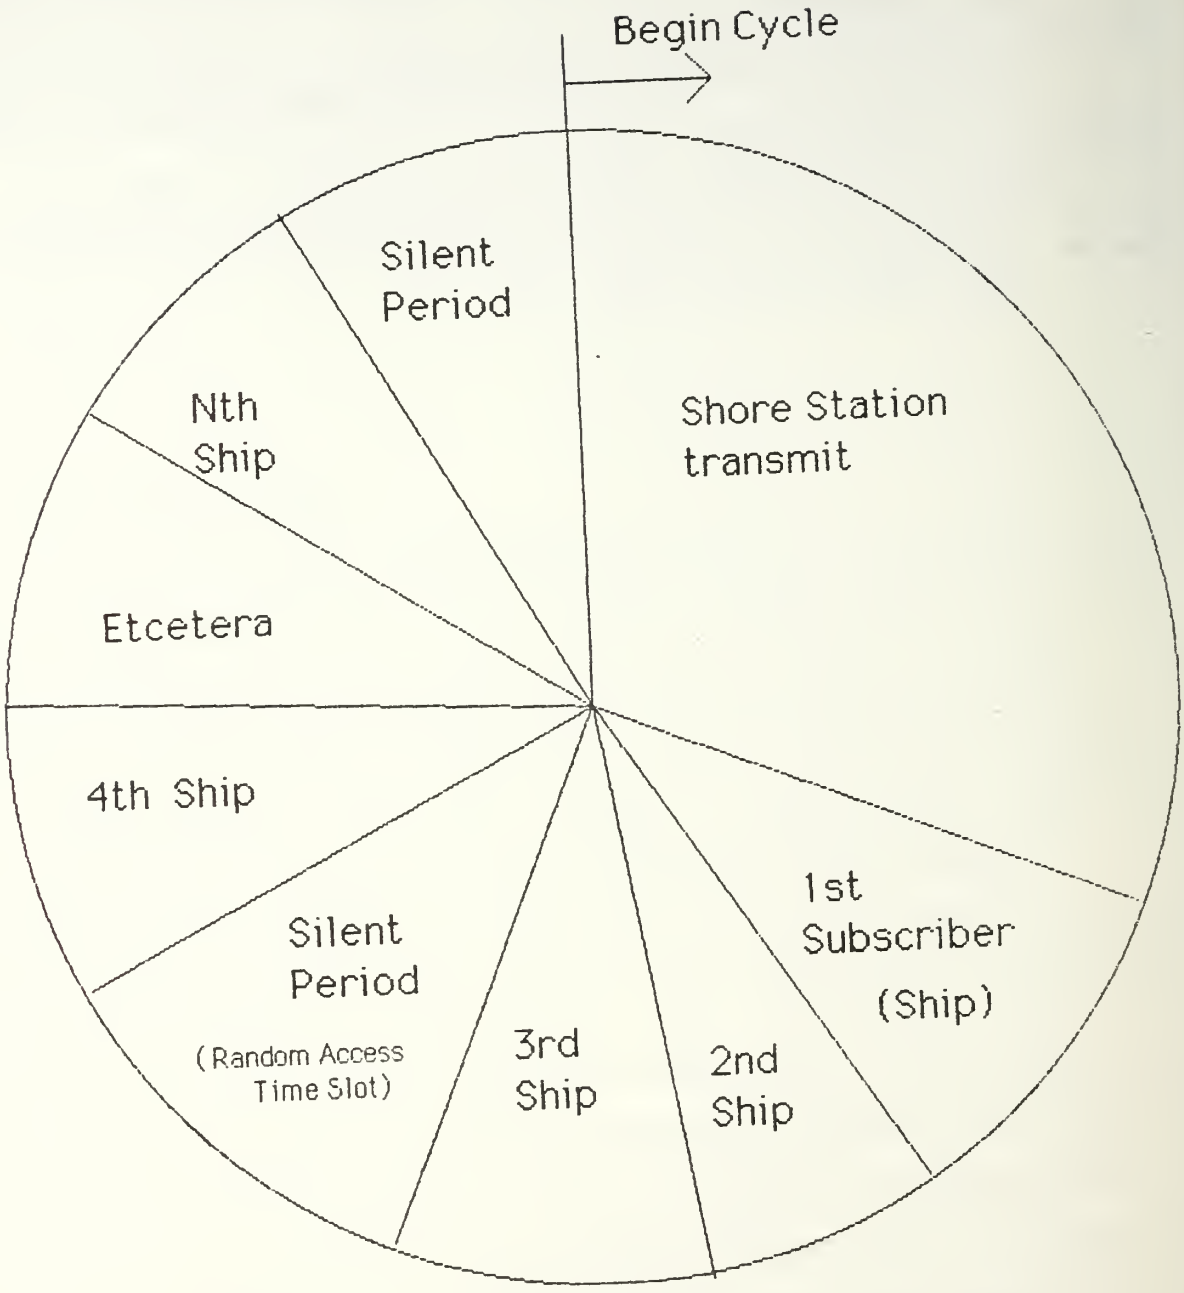

The CUDIXS Net Cycle

CUDIXS has some other features worth touching upon. Hardware is based on the AN/UYK-20 computer which is the military version of a PDP-8. This processing capability is currently used to the maximum by the existing software. There is no room for system growth.

The system is still essentially message based. Internally, however there is an ARQ scheme for handling errors. This error control scheme is adequate for the 1200 baud CUDIXS scheme where signal to noise ratios are quite good and the error rate is low. This difference in error rates also prohibits a direct importation of the CUDIXS architecture into the HF world. CUDIXS lacks and adaptive capability to cope with changing conditions.

One final problem uncovered in my investigation is the configuration control problem. Normally, the CUDIXS software is updated annually. The Naval Telecommunications Integration Center has taken great care to modularize the software suite so that most portions can be updated on a particular ship independently.

But some modules change the basic architecture and must be implemented simultaneously by all users across a communications area. Since the existing system cannot handle file transfers, the updates must be mailed to each command, and receipts obtained (assuring that each user has his copy), before the implementation can be placed into operation. For a fleet deployed worldwide, this is a major undertaking which requires several months to complete.

The lesson here is that configuration control is a critical aspect of the software engineering problem and all the lessons about modularity are very important.

F. NAVAL TACTICAL DATA SYSTEM--LINK 11

Link 11 is the automated system for passing tactical data among units of a battle force. The Naval Tactical Data System (NTDS) finds its roots in the anti-air warfare experiences of the Pacific Fleet in World War Two. When under attack by large numbers of aircraft, formations found that: 1) they were unable to fight the whole formation without extensive information sharing and 2) that manual (voice radio) means were rapidly becoming inadequate.

Out of this need, which has steadily become more critical, was born NTDS. The messages that are shared are short; thus individual logical messages (track reports) equated to physical messages (packets). Without ever quite realizing it, the Navy invented a LAN for tactical battle group use.

Network control in the first incarnation of NTDS was performed using a token ring type of architecture [Melich86]. Each ship was assigned a sequential number and each ship transmitted in turn. At the end of each transmission, the sending ship would 'sign off' thus signaling the next ship in the sequence to begin her transmission. This system worked until one ship missed its turn--then the system crashed. Because of the brittleness of this architecture, it was quickly abandoned in favor of a polled network architecture that is used today.

Today the net control station polls (roll call) each station in turn. If a ship misses its turn, the net control station recovers the net and signals the next ship to transmit. This revised polling system results in a more robust system--as long as the net control station remains operable. If net control can no longer operate, through failure, battle damage or jamming, an alternate net control must assume control of the net. In this respect, the system is still vulnerable, but this is the current tradeoff between efficiency and vulnerability.

Link 11 contains some error detection capability so that erroneous messages can be detected and discarded. Since track data is frequently refreshed, there is no need for a feedback loop (ACK/NAK) system to gain retransmissions of faulty messages. This is the classical example of the No_ack type of packet.

The original token ring uses the assumption that all ships in the network can hear each other, or at the very least each ship can hear the one sequenced before her. The centralized architecture assumes that all stations can, at a minimum, hear the net control station.

Link 11 operates in the UHF band for line of sight (LOS) ranges and in HF for extended line of sight (ELOS) ranges using groundwave propagation. NTDS is designed for intra-task-force communications, and is not directly applicable to the beyond line of sight (BLOS) system that this thesis is addressing, although the network accessing scheme is of interest.

G. COAST GUARD DATA LINK--SCAMP

This final survey subject is the Coast Guard's effort to develop a mechanized datalink that would operate over ship-shore links, including HF.

The genesis of this project lies in the procurement of the Coast Guard's 'standard terminal', an 8086-based microcomputer, procurement of which began approximately 1980. One of the standard terminal selection criteria was the desire to replace existing teletypewriter equipment with a service-wide standard equipment.

Simple equipment replacement can be performed with much less sophisticated equipment, so the sights of the developers are aimed a bit higher. The intent of the project, which is centered at Coast Guard Station Alexandria, is to provide service through ISO Level 2--the logical link layer.

Several characteristics of this project are similar to the thesis, but direction is significantly different:

- terminal equipment, the standard terminal, is specified. Whatever the software may consist of, it must fit into the already-specified hardware. Beyond that, this project assumes the same equipment as the thesis does: existing HF radio equipment.

- ISO Level 2 is an awkward place to partition the project. Many characteristics of the modem and cryptographic equipment have strong impacts at the network level. Without these considerations, equipment selected for data layer optimality may not be optimal across the entire system. Station Alexandria has recommended ADCCP (similar to HDLC) as the link level protocol. HDLC is the link level component protocol of X.25.

- dealing with networks will be done with a polling system. Beyond that, the polling system has not been described. A decision has not been made at the policy level, but the network level protocol will undoubtedly be X.25. As has been shown, X.25 has serious problems due to its foundation upon a full duplex physical layer.

The development effort is a bottom-up one, relying on the ADCCP and X.25 protocol specifications to prevent dead-end development.

APPENDIX D

SOFTWARE ENGINEERING

These two appendices (D and E) provide the basis for taking this thesis from a concept to a working communications system at sea. This appendix contains a brief introduction to software engineering for one not acquainted with the subject. The second is an acquisition plan in the context of the DoD procurement system.

A. THE SOFTWARE LIFE CYCLE MODEL

This thesis has been primarily concerned with building an architecture for a ship-shore communications system. With only a few exceptions, discussions of specific hardware has been deliberately avoided. Hopefully, since we now know what we want the system to do, we can make intelligent procurement decisions --specifications for protocols and both software and hardware.

The highest cost of the communications system described will be in the software development, not hardware. Indeed all the hardware required to implement the architecture currently exists and is readily available--there is no development risk. Therefore it is instructive to include the basic development models used by government and industry for software development.

Because the communications system development is primarily a software engineering project, rather than hardware, we can anticipate significant development costs at the front end of a project. In particular, development of specifications for and implementation of a Network Protocol is the critical segment and should require a competitive design acquisition strategy which entails paying multiple contractors.

The right half of this model is drawn from the work of Barry Boehm, currently collected in his book Software Engineering Economics [Boehm, 81]. Also see [Fairley, 85].

Comparison of Accepted Industry and Government Models

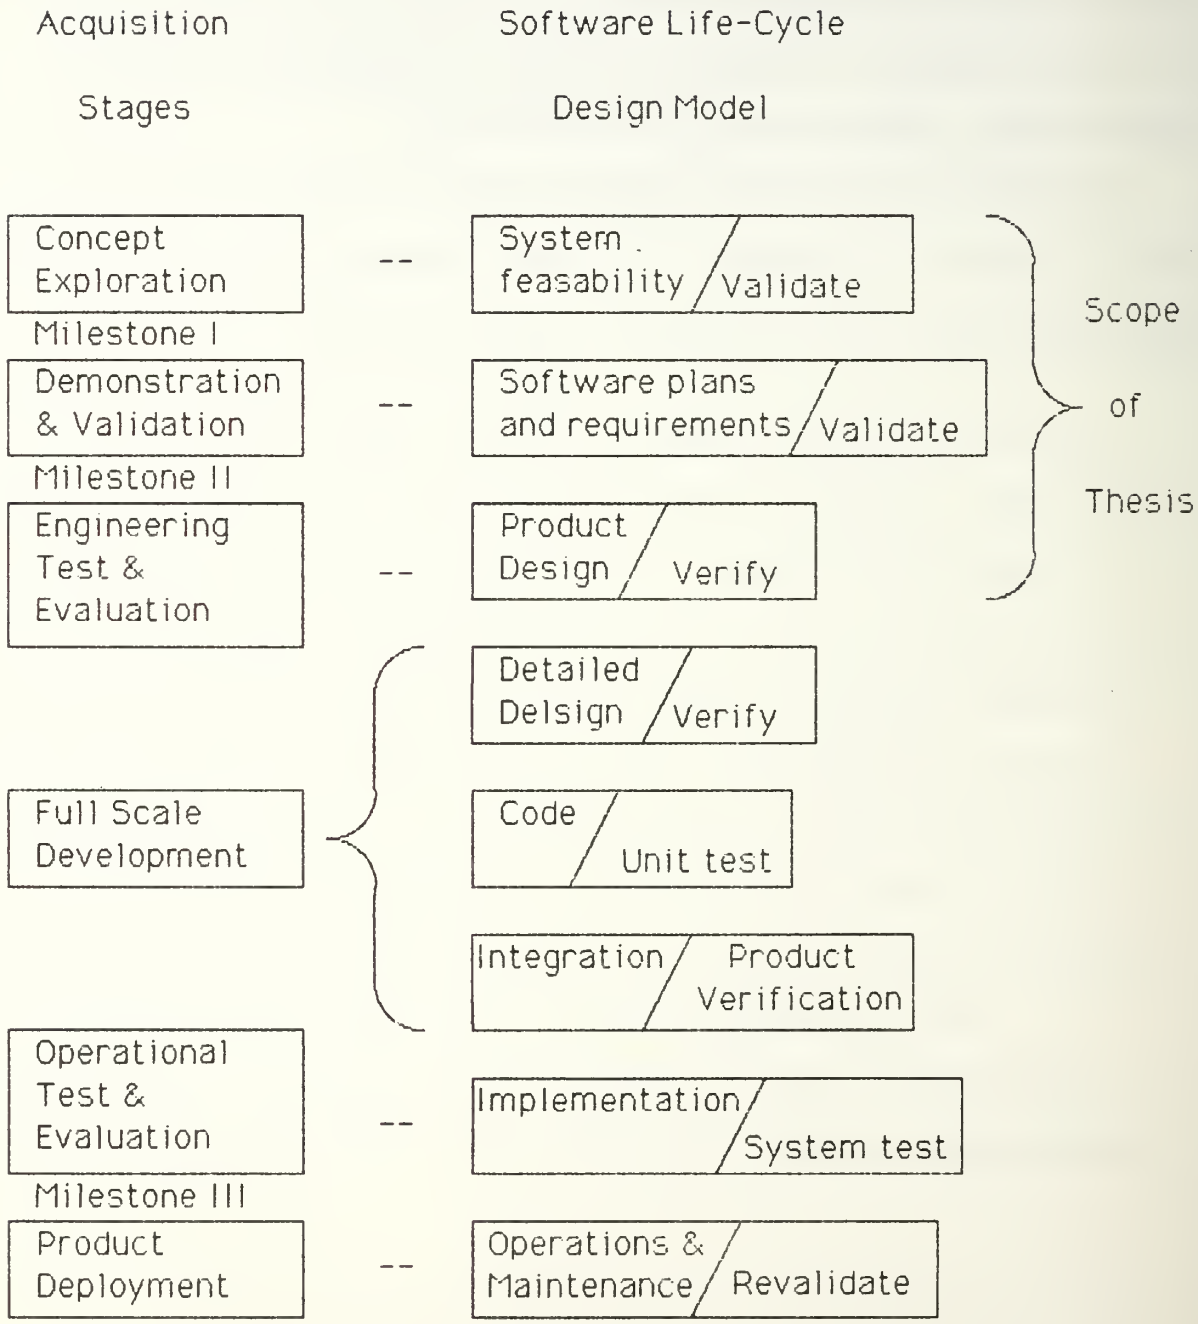

The left hand model is adapted from Navy Program Manager's Guide [NavMat, 85]. The resemblances between the software engineering and acquisition steps are so evident that it appears that the software engineers and the aqcuisition specialists are speaking parallel languages.

The primary lesson to be learned from software engineering is that one cannot get ahead of the game without paying for it in expensive redesign. For instance, attempting to perform product design (including hardware acquisition) before the system feasability and software plans and requirements phases are complete is wasting the taxpayers' money.

Said another way, if coding is attempted before the detailed design work is completed, this code will contain errors that will be very expensive to correct--usually resulting in the code being abandoned and the effort starting anew.

A third restatement, in the context of this thesis, is that if we attempt to implement a ship/shore system using existing inadequate standards without careful feasibility and product design stages we will end up with something that will likely be both inefficient and ineffective. This is painfully obvious when attempting to force the ship/shore communications needs into the X.25 protocol capabilities.

The process of concept exploration and iteration until validity is demonstrated cannot be shortchanged or forced without significant risks and costs later on.

This thesis attempts to rectify the frequently observed problem of program managers busily procuring equipment and specifying protocols without any comprehensive idea of the

overall system that these components are supposed to fit into. This thesis carries through the basic architecture--the first stages of the life cycle model.

B. THE ITERATIVE MODEL

Actually, there are two software design models, and they are not entirely exclusive of one another. The above depiction treats the basic design work as a one-time exercise. While that must be the case for the duration of a software or acquisition project, it will not be for the duration of sea service communications system life cycle. Therefore the second design model, an iterative one, is also useful.

The iterative idea both surrounds and is embedded in the life cycle model. Surrounding the development model is the process of using the lessons learned from operational usage in a redesign and updating process.

Within the life cycle model, each step must be passed sucessfully--if it isn't, then it must be iterated. This includes reexamining the assumptions passed from the previous stage.

Iterative models are usually referred to as prototyping. This school of thought finds its basic statement by Frederick P. Brooks who noted that the software designer had better "...plan to throw one away; you will anyhow." [Brooks75, p. 116]. Brooks' observation was that most programming efforts resulted in two codings of a problem. A prototype is one of the more accurate ways of estimating the software cost.

Once the initial program is installed, the iterative model is a software maintenance problem. Software maintenance is different from hardware maintenance in that software never wears out. Maintenance is concerned with correcting defects in the original program and improving the program.

Whenever software maintenance takes place, the second law of thermodynamics sets in. The maintained program generally loses some of the original structure and modularity--it gains entropy --unless the maintainer expends effort specifically targeted at countering this tendency.

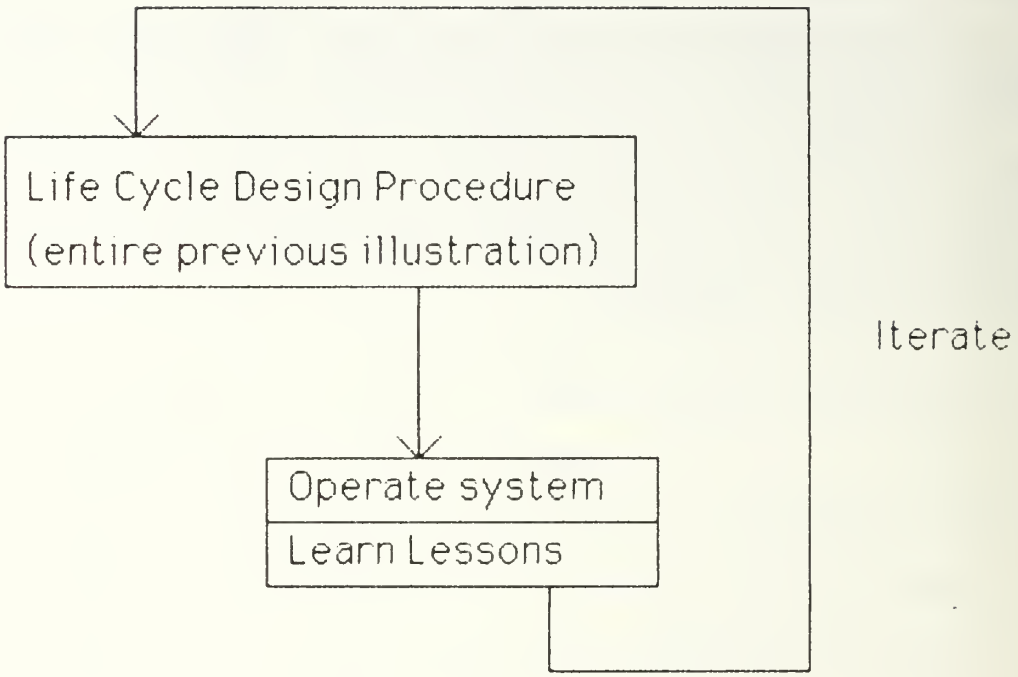

Iterative (prototyping) model of software development

Using the language of the military procurement community, the iterative model means use of a pre-planned product improvement idea. An analog in ship procurement is making adequate space, weight and power reservations so the shipboard systems have room for growth throughout the lifetime of the hull. The architecture of this communications system is intended to be open with room for growth.

This means modularity of development and a clear architectual conceptualization and room for growth. In this thesis, that has been done by careful specification of the boundary between nodes and links. Improvements in each area, and improvements in separate links (e.g. HF, VLF, satellite, shoreside links while in port) can proceed separately and concurrently without impacting the remainder of the system. And improvements in the ship-shore links can proceed separately without interfering with local area network issues within the ship or within the DDN ashore.

The third facet of this problem is that of configuration control. As an example, introducing communications-area-wide changes to the CUDIXS software has required several months and painstaking care in distribution so that part of the Navy does not find itself unable to communicate on changeover day. The difficulty in implementing changes fleet-wide weighs heavily towards keeping the number of changes to a minimum.

Because of the modularity, in particular the differentiation of links and nodes, it is practical to implement changes to nodes incrementally as ships return to home port. Most changes to logical link structures can be done incrementally; a few must be

done communications-area-wide simultaneously. But we now have an advantage in that changes to the HF links, for instance, will not affect the satellite system.

One footnote to the problem of software updates. Since the system described in this thesis is bit-transparent, it becomes possible to distribute the new software configurations on the communications system itself once the first configuration is installed.

C. SOFTWARE TESTING

Software testing is inductive. Testing can uncover errors, but it cannot deductively conclude that there are no errors in a program. Two types of testing exist:

1. Black Box Testing

This approach treats the program as a black box--the tester is unconcerned with what goes on inside the program, he is only concerned with the results.

To perform black box testing, the program is run and fed test data. There are several tricks that a tester can use to increase the rigorousness of the test including end point stressing, and heavy concurrent loading of a program that uses multitasking (as the node will). Nonetheless, the testing cannot be exhaustive, only thorough at best.

2. White Box Testing

This method of testing involves examining the code for errors. While this lacks the 'proof of the pudding' appeal of black box testing, it tends to expose the weaknesses in the program structure that make it difficult to maintain.

Both methods of testing should be employed. The white box approach is more useful early in the development process while black box testing becomes more useful at the production stages.

D. CONCLUSION

Software engineering exacts harsh penalties for shortchanging in the following areas.

1. Conceptualization

The technically perfect program is useless unless it performes the right function.

2. Top Down Planning

While it is possible to be too extreme, a systems view rather than a components view is necessary--and difficult. Lack of a 'vision' at the outset means that the end product will not have a cohesive function.

3. Modular Decomposition

Breaking the problem up into component parts requires great care in order that the system integration process work. Difficult as it is, however, the decomposition process is much easier than trying to assemble a complete structure from disparate components.

4. Planning for Change

Computer systems have chronically failed to have adequate room for growth. Many military systems are well known for their lack of modularity and extreme difficulty in maintenance.

As indicated by the discussion, these actions must be taken early in a project. Deferring the expenditures until later in the life cycle will result in much larger expenses and a less satisfactory product later.

APPENDIX E

TEST AND EVALUATION MASTER PLAN

A Test and Evaluation Master Plan serves a second function as the acquisition strategy document for smaller projects. This appendix is presented in that context.

A. MISSION.

Provide a modern, modular ship-shore communications architecture and system that is integral with the Defense Data Network (origin ARPANET) structure.

This system should be capable of efficient use of all customary ship-shore communications links in an integrated, stable manner. The initial effort will be bent toward HF links. The node specification--the Network Protocol which functions under the existing Transport Control Protocol/Internet Protocol (TCP/IP)--is general purpose and can handle any communication link that can be conceptualized as a one-way link.

Further, the Network Protocol can function alongside existing networks such as X.25. This allows for eventual interconnection of the ship-shore networks into the Defense Data Network.

B. KEY PARTS

A test and evaluation plan for the communications system outlined in this thesis should be broken down into the two customary parts: a Development Test and Evaluation and an Operational Test and Evaluation.

The chart on the following page illustrates the author's judgement of where development risk resides and how it should be managed.

Develop Network Protocol
Validate Network Concept
Build Prototype
Milestone I
Development Test & Eval of NODE
Network Protocol Specifications
Build Prototype LINK (HF)
Milestone II
Develop Production NODEs
Validate HF LINK
Develop Other LINKs :
- LF, VLF
- Phone
- Satellite
- Shipboard
Validate
Milestone III
Produce
Produce
Acquisition Strategy

C. STAGE ONE--PRE-MILESTONE I

The purpose of this stage is to design and validate a Network Protocol along the lines described in the thesis. Since the Transport Control Protocol and Internet Protocol are both available, they should be used without change.

The results of this stage should include:

1) Validity judgement of the network acknowledgement concept. Does the Network Protocol idea work?

2) A Network Protocol specification. This would be a refinement of Chapter Three of this thesis.

3) A software validation suite. The Network Protocol can be tested by simulating various links with software stubs. Similarly, the prototyped Network Protocol can be used as a driver to test prototyped links at later stages.

The Network Protocol specification need not be final in every detail. What should be stabilized, however, is the interface definition between nodes and links (Network and Logical Link layers). This is important because this definition allows work to proceed concurrently at minimal risk.

A prototype of the Network Protocol should be installed on a computer under an existing TCP/IP environment and tested. Software test drivers simulating data and control packets--both correct and flawed--from links can be used for this stage of the test. These test drivers should be fully documented as they form part of a test suite.

1. Critical Test and Evaluation Issues

Since this plan is written by the same author as the basic thesis, the reader should beware of the likelihood of biases and blind spots. In my subjective opinion, the likely places that may fail in this conceptual test include:

1) Inability to effectively handle different types of data (e.g. Full_ARQ, NAK_only and no_ack). In particular, NAK_only packets require error control at the transport level to detect missing packets which may not fully meet needs.

2) Inability to integrate packets from several links into composite messages. The system should handle different length packets, overlapped packets, etc.

3) Duplicate packets. Because of the redundancy needs of sea service communications, packets/messages will arrive multiple times on multiple channels. The node must be fully capable of handling large numbers of duplicates.

4) EMCON management. Two nodes must be able to operate effectively when one is unable to transmit. This means that the sending node will accumulate large amounts of un-acknowledged packets while waiting for the non-transmitting node to emerge from EMCON. The system must automatically account for entry into EMCON and again automatically clear the backlog when the non-transmitting condition is cleared.

5) Multiple addressing. A packet addressed to several ships should remain in the wait queue until all ships have acknowledged it.

Items 1) and 2) are not vital to a prototype if they are accounted for. If the problems exist, but reasonable assurances also exist that they can be corrected, we can safely pass Milestone I. The important issue is validation--is the concept correct? Verification--has the Network Protocol been programmed correctly--is clearly of secondary importance at this stage.

The DoD standard Transport Control Protocol and Internet Protocol (TCP/IP) should be used. If deficiencies in these protocols are uncovered, they should be documented as well.

D. STAGE TWO--DEVELOPMENT TEST AND EVALUATION

This stage involves building and operating a complete prototype of a node and link combination. Because the validity

of the network acknowledgement concept must be demonstrated to pass Milestone I, we are justified in proceeding with node and link development concurrently.

Step 1) Build some limited production nodes. This should mean little more than duplication and installation of the prototype software developed in Stage 1 on existing computers.

The node is certainly a candidate for software maintenance -- updating of the prototype software. We can safely assume that at least one complete rewrite of the Network Protocol implementation will be needed sometime before production. This need must be planned for.

Step 2) Build a few links. Here we should aim for the link with high potential payoff--HF. HF links are needed by both the Navy and the Coast Guard. Additionally, prototypes of HF links can be used to validate the network access protocol.

Deliverables at this stage should include:

1) Limited production of nodes (perhaps a dozen) with documentation suitable for operational (shipboard) testing.

2) Link termination equipment enough to build and operate a network in one communication area. These links must also be suitable for operational testing.

3) Preliminary operation and maintenance manuals for the communications officer, radioman, and electronics technician.

4) Documentation for the software maintainer.

1. Critical Test and Evaluation Issues

In addition to reviewing the Stage 1 issues, the following become important:

- software implementation. Concept validation was a Stage 1 issue, now the issue is implementation verification. Modularity of software programs and consequent maintainability is important. So is adequate documentation. A recommended Network Protocol standard should be ready for adoption.

- The only critical area of an HF link is the controller. A prototype should demonstrate the ability to handle the following situations:

    o. Multiple senders ganged together in a downward multiplexed configuration.

    o. Adaptive error correction capability. (Actual operation is secondary to a demonstrated ability to incorporate this feature at a later date.)

    o. EMCON controllability.

- Electromagnetic compatibility. In building links, the developers must begin to be conscious of RFI/EMI considerations in assembling hardware. At this stage, the consideration is preliminary as specific hardware is not yet identified. But habits such as avoiding switching power supplies, correct cable routing and proper bonding of equipment should be developed since production contractors are likely to use the prototype equipment as the basis for their manufacture.

    Self generated noise, particularly from digital equipment and supporting power supplies, is reasonably easy to 'plan' out of a system, but somewhat more difficult to fix after hardware is installed.

- required EMP protection. Whether this is to be accomplished by hardening equipment or providing redundant equipment must be addressed.

- communications operator and manager interface. Supervisors must be able to isolate and repair failures easily and without disrupting the complete network.

At the conclusion of this work, the project should pass Milestone II.

E. FULL SCALE DEVELOPMENT

At this level of testing, it is essential to get the systems aboard ships and get the ships to sea. Communications in

general, and HF in particular, tend to show up lab tests as unreasonably optimistic.

Operational Test and Evaluation. The new system should be tested in parallel with existing operating systems. The testing should include:

1) high traffic situations. Both in terms of multiple links to a single ship and multiple ships on a network. This is intended to validate both the integration and network access concepts.

2) operation in high noise, or poor HF propagation situations. This should validate the logical link layer work. It may be suitable to test in a jamming environment, but this is dependent on the physical equipment attached.

The manuals should be revised based on feedback from these tests. Training and logistics plans (maintenance philosophy) should be generated at this point.

F. MILESTONE III REVIEW

At this point, the software needs will be pretty well defined. In preparation for Milestone III, target computers for nodes and links should be identified. The software must be ported to hardware that meets the physical needs of the platforms that carry it. These needs include making a version of the system physically small enough to install in patrol boats and aircraft. In addition to the usual EMI/RFI, EMP, vibration, and survivability issues, the computers should be sized to permit growth:

- as the software is maintained, it can be expected to expand and require more memory. Either a paged memory operating system or adequate RAM to permit at least 100% anticipated growth over the prototype should be planned.

- since unacknowledged packets will probably be treated as multi-tasking operating system processes within the Network Protocol implementation, the ability to handle a large number is required.

Three sizes of nodes are probably required. The largest is a communication station node which should be capable of supporting a large number of links and users. The second should be sized for major combatants and auxiliaries (FF and up). The third should support small units with one or two each senders and receivers. Note that the differences are in physical hardware, not the software.

At this point, the communications system is ready to clear Milestone III and proceed to production.

G. PRODUCTION TEST

At this stage, we are ready and capable of producing specifications to procure production nodes and links ready for installation. As the contracts are generated, the following items should be considered:

- revalidate concept and software--review specifications and protocols.
- build in the remaining production requirements such as backup power supplies, installed spare equipment etc.
- allow for addition of new links and updates of existing links. Ensure adequate growth reservations.
- ensure EMI/RFI and EMP immunity adequate to the need.

A maintenance plan should be generated at this point. It should be made up of two parts:

1. Hardware Maintenance

This is the customary problem of providing adequate spare parts in inventory to handle breakage. The questions of user maintenance versus depot turn-in must be resolved here.

2. Software Maintenance

Software maintenance is fundamentally different from hardware maintenance because software does not wear out. Maintenance occurs for two reasons:

a. Defects (bugs) appear through use.

Hopefully, the development process has reduced both the number and the magnitude of defects, but it cannot eliminate them entirely. Software testing can only verify the presence of defects, never their absence. No matter how careful the implementation, some bugs will appear and the software maintenance system must be able to cope with them.

b. Enhancement

As the sea services gain experience, the need or desire for improvement will become evident. As the specifics of these desires become apparent, the software maintenance effort should respond. It will be at this point that 1) inadequate reservations for system growth and 2) poor initial software engineering and modularity--will appear. Money, time, and effort spent at the design stages will here be repaid many times over.

APPENDIX F

A COMMUNICATIONS LINK FOR THE MARITIME DEFENSE ZONE: A LORAN STATION REBROADCAST

This communications link uses Coast Guard LORAN stations as senders. It is a derivative of the Clarinet Pilgrim system which used LORAN stations to copy a Navy fleet broadcast and retransmit it by modulating the signal onto the LORAN signal.

A. LORAN STATIONS

The LORAN-C carrier frequency is 100kHz with a pulse modulation. The Low Frequency (LF) band is used because the ground wave signal can be propagated far enough for navigational use. Sky wave signals, with their varying paths, and consequent degraded accuracy are not used for navigation (though they exist in this band).

LORAN signals are broadcast by each station in pulse groups. Each group has eight individual pulses (master stations have a ninth pulse for identification). Depending on the chain, the pulse group rate varies from 25.00 pulse groups per second (40000 microsecond group repetition interval) to 10.01 pulse groups per second (99990 microsecond interval). Each station in the chain broadcasts at the same rate.

Current LORAN-C navigational coverage covers at least 100 miles offshore throughout the area known as the Coastal Confluence Zone. Navigational coverage means that a receiver can receive at least three signals--from the master and two secondaries--within geometric fix accuracy limits. For our

purposes, it is important to note that signals can be received well in excess of the navigational limits of the system.

Advertised LORAN coverage is rarely bounded by signal strength, but rather by chain geometry. At the point where the hyperbolic lines of position intersect with angles of less than 15 degrees, accuracy is considered to be impaired. Navigational quality coverage encompasses virtually all of the United States Fisheries Conservation Zone (geographically identical to the declared Exclusive Economic Zone). This also includes all of the continental shelf areas where the Maritime Defense Zone commanders can be expected to conduct operations. Also to be noted, LORAN coverage is not solely maritime--most of the continental United States and Canada is fully covered by navigational quality ground wave.

Two new chains are currently in the planning stage; one is an Alaskan North Slope chain; the second, a Mid-Continent chain. With these two chains on air, full coverage across the entire United States landmass (and most of Canada) can be expected. The reason for these chains is to support FAA navigational requirements.

B. USING LORAN TO COMMUNICATE

The communications link operates by modulating (using phase shift) a data signal onto the LORAN-C carrier. Any receiver that can receive the LORAN signal receives the communications modulation. The communications signal can be modulated onto any or all of the eight pulses the LORAN signal. Each pulse can carry one bit of information.

Various systems have actually been built in the past using one, two or six of the pulses for communications. Thus there is no technical risk in such a system.

For communications purposes, only one station need be received--the coverage is usually much wider for communications than navigation. This covers all Coast Guard operating areas except Antarctic icebreaking.

To implement LORAN station broadcasts within the context of this thesis, we need only treat a LORAN station as a sender. A data stream from the communication station is fed to the LORAN station. The only remaining technical function is to modulate the data stream from the communications station onto the LORAN carrier. Assuming that the signal is encrypted at the communications station, no security interface exists at the LORAN Station so there is no classified liability at the LORAN station.

The following illustration diagrams a sender for LORAN station broadcasts. The basic structure is similar to that of an HF sender, but several parts--especially the cryptographic process--are physically located at the communications station rather than the LORAN station. This relieves the LORAN station of the necessity of holding any classified material.

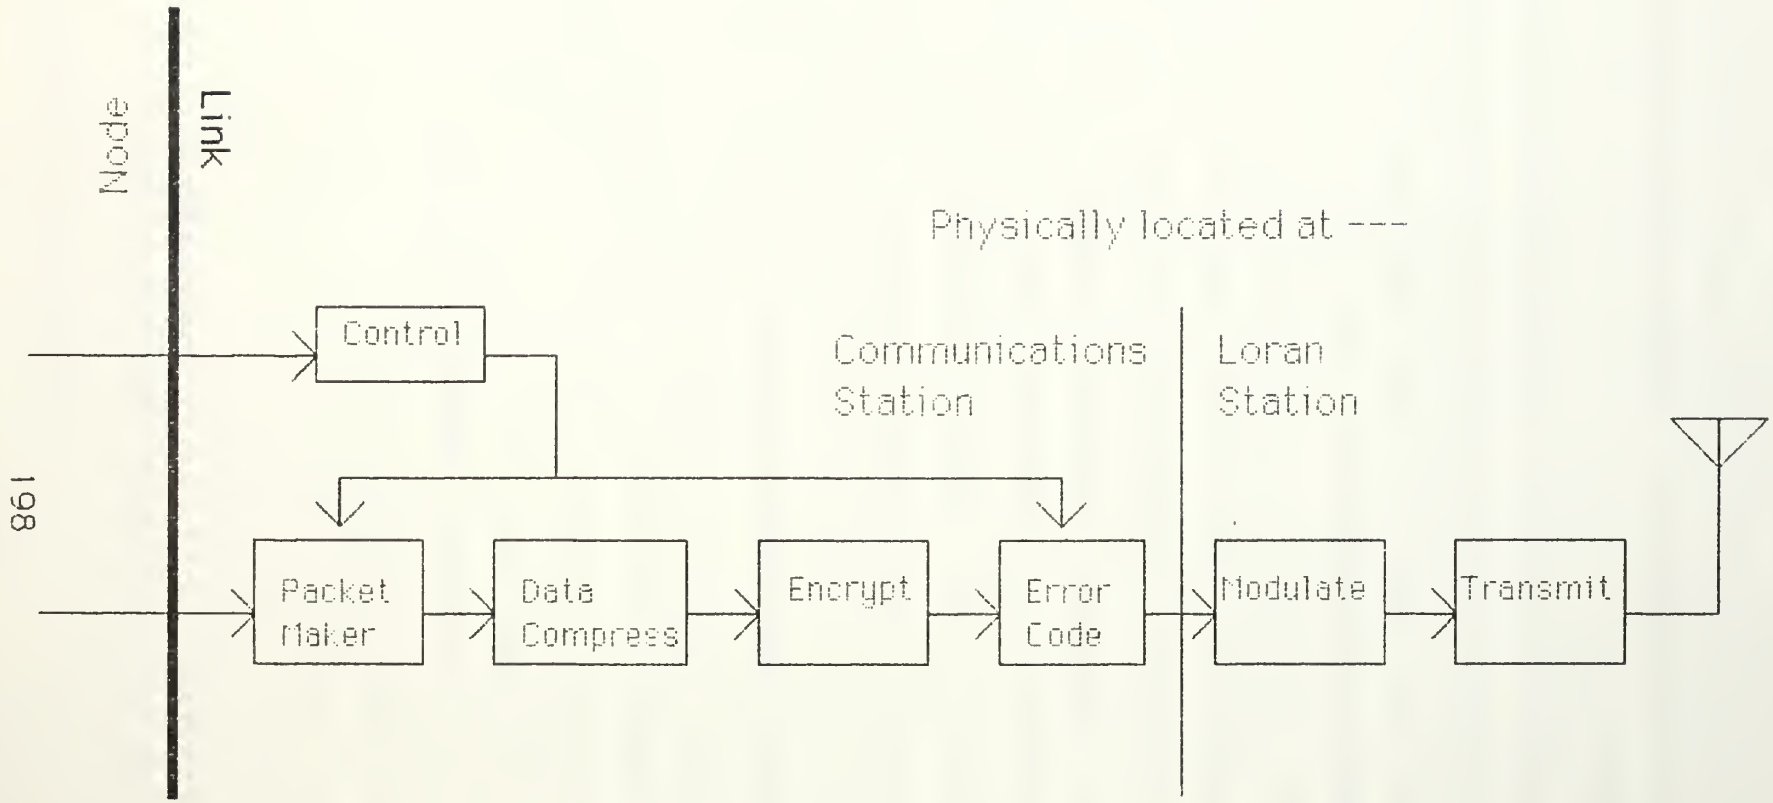

Sender -- Using Loran Transmission

For the receiver portion, monitor receivers currently used at LORAN transmitting and monitor stations are capable of interpreting the modulation and can be easily adapted for shipboard use. A block diagram of a receiver would be essentially identical to the HF receiver earlier described.

Because of the wide coverage and because of the high on-air requirements (99.7% on air, in tolerance) of LORAN stations, this system has potential emergency national communications implications beyond its use to the Coast Guard. All LORAN stations have backup power supplies capable of operating the station in the event of commercial power failure. There are currently 23 transmitting stations in the United States (including Alaska and Hawaii) and 3 in Canada. Because of their highly distributed nature (nobody can find Searchlight, Nevada or Baudette, Minnesota even when looking for them), the probability that some stations would survive even a nuclear attack is pretty good. The difficulty would be getting the data stream to them to transmit.

1. Speed

LORAN pulses are identified by their Group Repetition Interval (GRI). For instance, the West Coast Chain, which covers the Monterey area, has a GRI of 9940, meaning a 99400 microsecond interval between transmission of one pulse and the next. This works out to just over 10 pulse groups per second. If all eight pulses in a group are used for communications, the data rate is 80 baud.

If the baselines in a chain are shorter, such as the Central Pacific Chain (Hawaii), the GRI is decreased--in this case to

4990, which increases the baud rate to 160. For navigational reasons, it may be prudent not to change the modulation of all of the bits in a pulse. If, for instance, the first two bits were not used for communications, then the baud rate would decrease by 25%.

2. Configuration

There are two ways to configure such a system, neither of which are mutually exclusive and which have tradeoffs.

The first is to feed a data stream to each LORAN station independently, as shown in the following illustration. Since each LORAN station functions as a sender in isolation of the other LORAN stations, the total throughput is the sum of the maximum rates of each station. Three stations, each with an 80 baud capability, can provide a total of 240 baud.

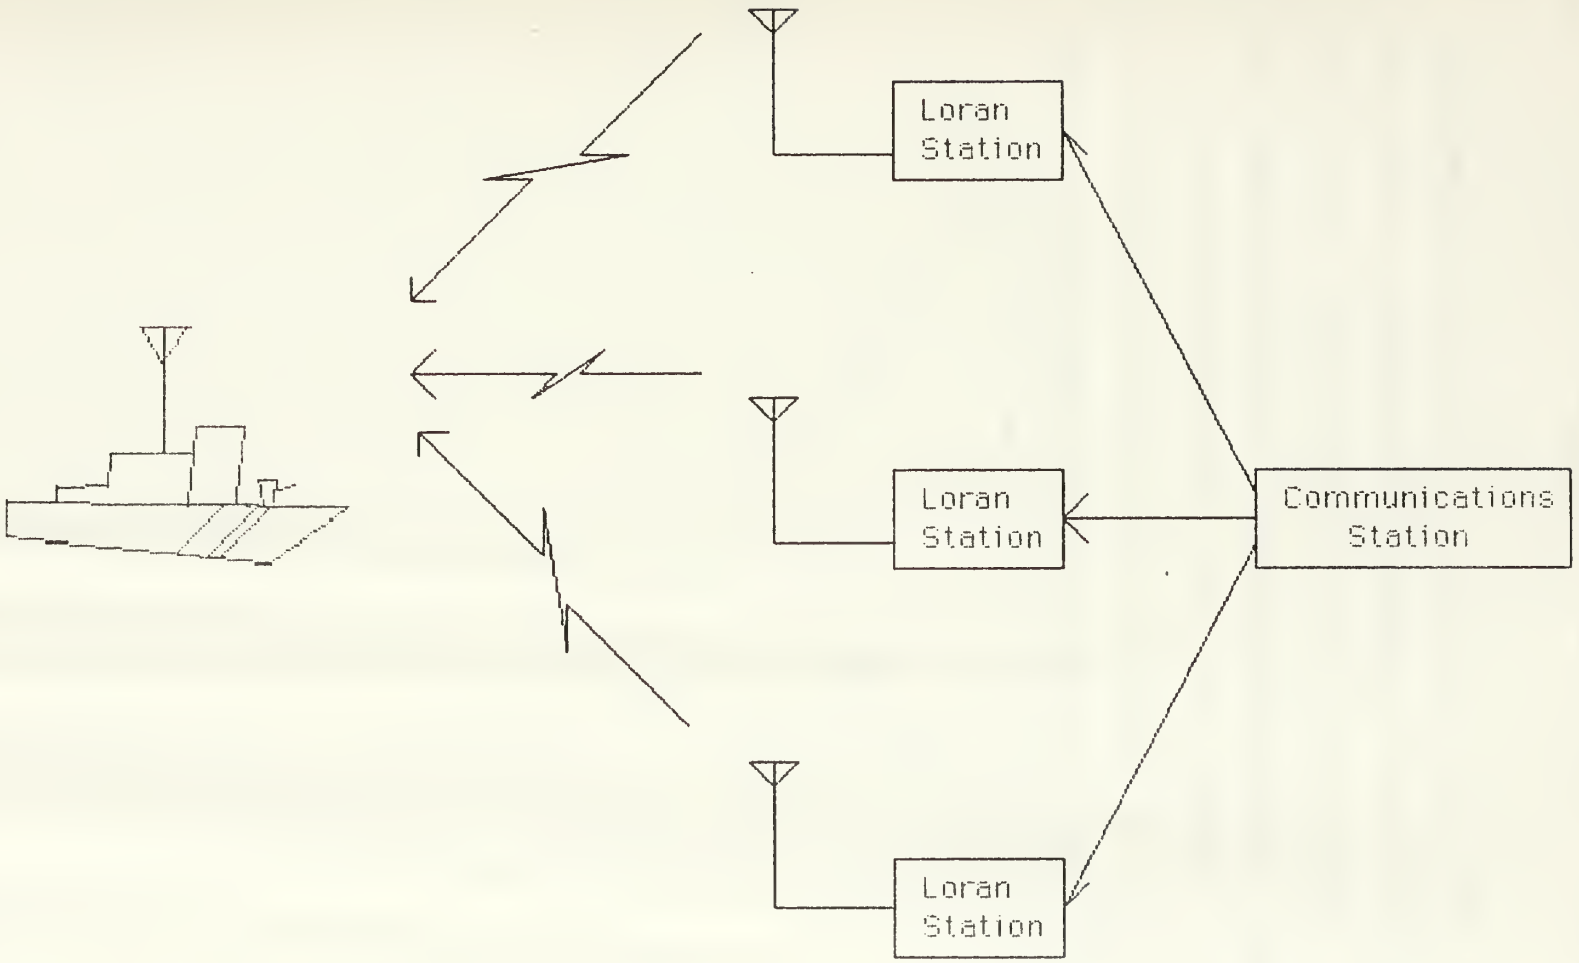

Higher Data Rate -- Lower Redundancy
Configuration

The second configuration sacrifices the higher data rate to improve redundancy. Each station modulates the same signal. Thus a receiver that can receive any of the stations can copy the data.

This allows some shortcuts in shoreside connectivity. It is only necessary for the communications station to feed one LORAN station. The rest can copy the first signal and rebroadcast it. This allows a single-thread connectivity arrangement that is less expensive to install, but more vulnerable to failures. It should be noted, that if the distributed system is installed, it can be reconfigured easily to the single-threaded system if some shoreside links fail--but the reverse is not true.

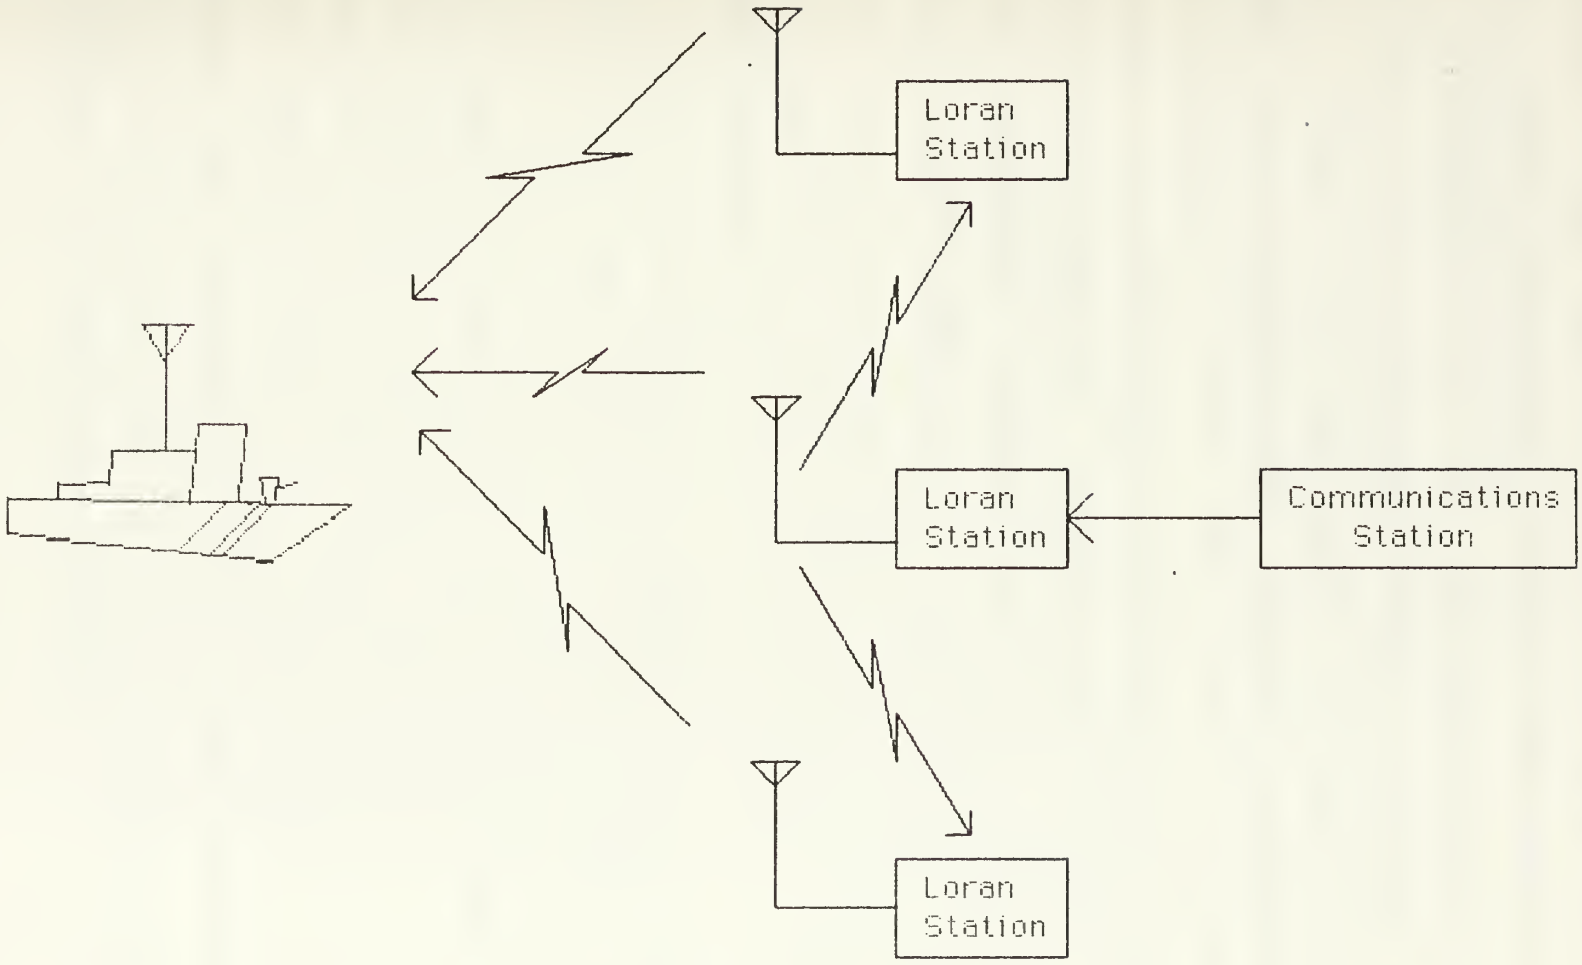

High Redundancy -- Low Data Rate
Configuration

C. CONCLUSION

Using the LORAN system as a Maritime Defense Zone shore-to-ship communications link is highly feasible. The developmental costs have already been expended in developing Clarinet Pilgrim --there is no development risk. The costs of constructing and operating LORAN station transmitters and receivers is sunk--the communications capability requires only marginal costs.

Additionally, a latter day Clarinet Pilgrim link is somewhat easier to design than the HF link. Since the navigational timing requirements already keep separate transmitters from overlapping, there is no network access problem. This makes the controller simpler to design. The LORAN rebroadcast technique is fully integrable into the context of one way links used in the thesis.

APPENDIX G

SECURITY CONSIDERATIONS

A. AN ENCRYPTION CONSIDERATION

Link encryption using a machine cypher system is entirely practical--the KG-84 has more than adequate flexibility to easily integrate it into our system. Conceptualizing links as one way channels considerably eases the integration of data stream equipment.

There is one caveat that this writer feels must be made clear to the cryptographer, however. It is this.

We now know that the Allies had thoroughly broken the German Enigma code and the Japanese Naval codes throughout World War II. We also know that the biggest chink in the cryptographic armor was the recurrence of certain character patterns in every message [Kozaczuk, 79].

Whether this vulnerability remains in our cryptosystems is beyond both the scope, classification and knowledge of this author and thesis, but it should be made completely clear that packet systems, by their nature, have recurring character patterns throughout. For instance, the first byte in a packet is a flag byte (01111110 in X.25) pattern, as is the last. Other bit and character patterns, such as packet headers, recur predictably and repeatedly. If the cryptographic algorithm of the KG-84 system allows these patterns to jeopardize the security of the system, it should not be employed.

B. SHIPBOARD INTERFACING

In Chapter Three, an illustration of several one-way links, and one conventional network appears. This could be an example of a shipboard node where the one way links are the ship-shore circuits and the conventional network is a shipboard LAN designed to connect everybody from Combat Information Center to the engineering supply petty officer together. While the architecture can now be fully envisioned, there is a security problem that must be solved before such a network becomes reality.

The communications system must reliably guard against the possibility of a classified message appearing on an unclassified terminal. There are three approaches to the problem.

1. Provably Correcty Software

Require that the nodal software--the Transport Control Protocol, Internet Protocol and Network Protocol--all be provably correct. This approach has two major drawbacks:

a. Provably correct programs are extremely difficult to realize, very expensive and equally difficult and expensive to maintain.

b. Errors introduced into the data stream by the communications medium may result in mis-delivery despite the correctness of the software. This can be minimized by the header checksum, but packet correctness cannot be absolutely guaranteed.

Provably correct programs is not a practical solution.

2. System High

Run the communications network as a System High one. That is all terminals be cleared to the highest classification of the system. This is how the current ship-shore system operates.

The drawback is that the people needing only unclassified access must be fully cleared, not because they need to see classified material, but to operate the system. For a system with limited terminals, such as the existing one, that is a poor, but workable solution.

The System High approach is increasingly less practical. As the system supports more terminals and users, this becomes impractical. Operators all along the data path see messages that they have no need to know the contents of. This allows compromise at any node to compromise the whole system.

3. Use end-to-end Encryption

Since the communications system outlined in this thesis is bit-transparent, it has no need to know the content of the bits that it is sending--if those bits happen to be encrypted before they are committed to the communications system, that makes no difference.

This denies neither the need nor the capability to do link encryption. But full use of end-to-end encryption for all classified traffic means that the primary purpose of link encryption becomes resistance to traffic analysis rather than protection of the traffic itself.

In a shipboard network, incoming traffic gets routed to various terminals. Each terminal has an end-to-end encryption device with a key suitable for that terminal's level of classification. Consequently, if a secret message gets accidentally routed to an unclassified terminal, it cannot be decrypted and read.

APPENDIX H

COMMUNICATIONS SYSTEM DISTRIBUTION

This appendix carries the downward multiplexing notion of Chapter Two one step further. This is a follow-on consideration that can be considered more fully once an operating system is built.

A. COMMUNICATIONS STATION DISTRIBUTION

Currently, the US Navy depends on four Communications Area Master Stations (NavCAMS) and one large communications station (Stockton) to carry the bulk of the communications traffic load.

This concentration of resources has significant savings in resources in terms of manpower and equipment savings--economies of scale given the existing architecture are significant. The drawback is that this concentration is dangerously vulnerable, assuming that the communications stations will be targeted in a general conflict. This concentration of resources also increases vulnerability to sabotage short of war.

With the architecture of this thesis, we can consider decentralizing the communications station function; this appendix describes the potential.

Up to now, the thesis assumes that the communications station distribution will be similar to the existing architecture. The station itself is a node and a series of terminating equipments are required, one to a link.

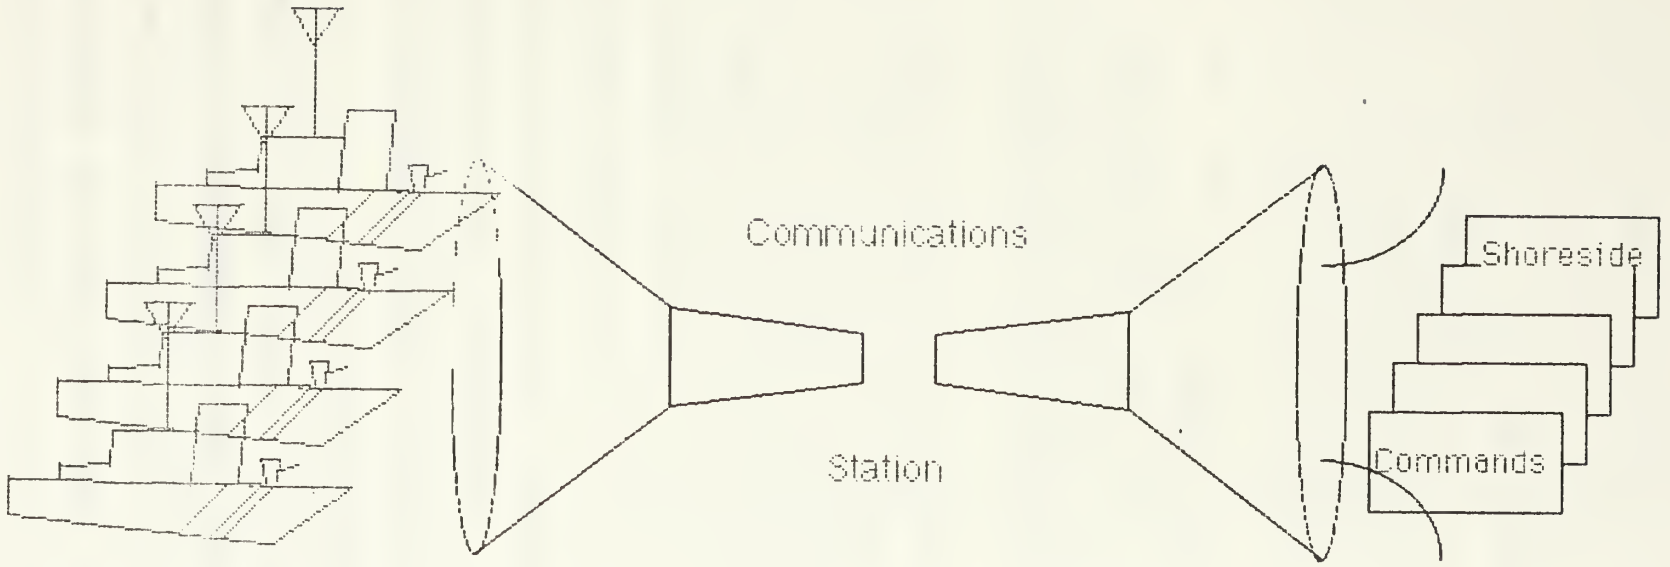

Centralized Communications Station

The currently installed node equipment is the Naval Communications Processing and Routing System (NavComPaRS) computer. NavComPaRS predates the ARPANET development work so it doesn't map to the ISO model well, but it is performs network layer functions. Given the state of technology at the time of its implementation, there were significant economies of scale to buying one mainframe computer to serve a number of terminating equipments.

Two developments have changed this:

1) the development of dirt cheap computing power,
2) the development of network technology and advent of the Defense Data Network (DDN) as an operational communications system.

Visualize the communications station function, not as a mainframe, multiuser computer, but as a network of nodes. Failure of any one node does not affect the rest. It becomes very feasible to relocate the termination equipment anywhere a node can be plugged in to DDN. Practically speaking, any Coast Guard or Navy shore unit with a DDN tie--even a node with a commercial telephone dial-up TAC connection--can become a communications station. Furthermore, the equipment can be quite mobile. From an enemy targeting view, this makes it much harder to degrade the communications system.

Today, a well aimed hit can completely knock out a communications station and leave a communications area largely bereft of service. With a networked system such as the one illustrated on the next page, a hit can only damage one node and the links to it. Vulnerability can be reduced by procuring more equipment and distributing it.

Our downward multiplexing model showed that it is practical to set up multiple links between two nodes (communications staion and ship). It is just one more step to physical separate the shoreside nodes from each other.

To properly perform the distribution function, the Network Protocol needs to function across several nodes so one node can support a series of senders and another, some distance away, can support receivers. This way, transmitters and receivers can be located far enough away from each other to avoid the full duplex inhibitions that the ships must face. Control packets that arrive at the receiver site must be passed to the transmitter site because that is where they are used.

We will not fully develop the notion of a distributed shoreside half of the ship-shore communications system. Rather, understand that the architecture illustrated offers the possiblility in the future.

B. SHIPBOARD DISTRIBUTION

Similarly, a battle group can also distribute its communications. For instance, the carrier may receive on HF but beam its acknowledgement stream to an outlying ship (or aircraft) by line of sight frequencies. The outlying ship then transmits the data ashore by HF. This avoids pinpointing the carrier.

Variations of this technique have been used in the past (one arrangement was known as AUTOCAT) but full networking of the battle group has not been very practical within the existing communications architecture.

APPENDIX I

DATA COMPRESSON

Data compression allows us to get more information through a given amount of bandwidth. For example, microcomputer communications have routinely used the SQueeze/UnSQueeze programs to cut down on telephone connect time.

This appendix discusses criteria and use of data compression mechanisms.

In function, there are several compression algorithms designed for different purposes. The SQueeze/UnSQueeze mentioned above use an algorithm known as Huffman coding. A squeeze program first counts characters in a message and creates a statistical frequency table. The most frequent letters (usually spaces and 'e's) are then coded with the shortest bit forms. Less frequent letters such as 'q' and 'z' are given longer bit forms. Then the file is compressed by substituting the bit forms for the characters. Huffman coding generally yields a 1.8:1 to 2:1 compression efficiency.

One other form of compression is Lempel-Ziv encoding. This system uses repetition of characters to drive the algorithm. Because Lempel-Ziv encoding requires a minimum file size of around 1000 bytes before it becomes effective, it is of little use as a link encoding algorithm and is not discussed further.

Tokenization is a computerized version of brevity codes that the military has always used. Words and phrases are replaced by tokens that are shorter. One commercially available tokenization scheme indicates efficiencies of 3:1 to 4:1 for text files, making it quite attractive for our system.

In form, there are two compression implementations: dynamic and static. The conventional Huffman systems are dynamic. This means that the coding scheme--the statistical frequency table --is included within the file transmitted. The drawback is that if the coding scheme is damaged, the file will not fully decompress. Errors will propagate throughout the data compressed.

Static systems embed the compression/decompression scheme within the compressor itself. This decreases the performance of systems like Huffman coding, but tokenization schemes can remain quite effective providing the token library is sized properly and optimized to look for standard phraseology and frequent bit patterns.

For robustness reasons, static systems are best for our HF system (or any system expected to operate in a military environment). An additional reason for tokenization schemes is that they can effectively operate at the logical link level on bit streams passing through the compressor, where Huffman systems conventionally operate on a complete file in a batch.

The payoff. If a compression efficiency of 3:1 can be achieved, then the effective information transmission rate is tripled over the raw baud rate. Thus, if our circuit is operating at 2400 baud, it will appear to be transferring 7200 baud worth of data.

The pitfalls. If the data sent is end-to-end encrypted, it will be presented to the compressor as a string of random bits. Excepting the packet header, the tokenizer will not be effective. Similarly, if the data presented is not textual--containing recognizable character strings--no compression can take place. This will also be the case for graphic data and computer programs (although database transactions may have recognizable phrases). Thus the efficiencies of tokenization are fully realizable only for textual data. It should be noted that even if no compression takes place, the compressor will not damage the data by being in the system. The worst that can happen is no gain.

An effective complementary alternative to link compression is end-to-end compression. Similar to end-to-end encryption, this means compression at the application level where a compression algorithm suitable for the particular data can be effectively used.

The two approaches, link compression and end-to-end compression are complementary, just as link encryption and end-to-end encryption are. Our system is able to take advantage of the benefits of both.

In order to do this, a link compressor, probably the tokenizer described above, should be used in the system at the logical link level. This compressor should be optimized to

compress English text, specifically the verbiage used in sea service messages.

End-to-end compression should be employed by specific systems. For instance, the communications manager should require the fleet oceanographer to implement end-to-end data efficiencies in his weather reporting and forecasting systems. An example would be a compression algorithm in the FAX map digitizers and plotters.

Although the exact design is not critical to the architecture study of this thesis, a tokenizing data compressor was envisioned in the link senders and receivers.

ACK. Acknowledgement.
An acknowledgement of receipt in computer communications terminology. Derived from the ASCII character 06H, ^F. See QSL.

ACRONYM. A Collected Rabble Of Nonsensical terms Yielding no Meaning.

AIG. Address Indicating Group.
An address that is itself a list of addresses. This relieves a message drafter of the burden of remembering all the addresses that a particular multiple address message is to go to.

ARPANET. [Defense] Advanced Research Projects Agency NETwork.
Experimental/developmental packet switched data communications network. See DDN.

ARQ. Automatic Repeat reQuest.
A communications procedure where a unit of traffic is retransmitted until received successfully. See backward error correction, ACK, NAK, QSL, ZDK.

ASC. Autodin Switching Center.

AUTODIN. Automatic Digital Information Network.
AUTODIN and AUTODIN II are existing message switched Defense Communications System backbone carriers.

AX.25.
The amateur radio variant of the X.25 standard.

BER. Bit Error Rate.
Digital measure of errors introduced by the transmission medium in a communications channel. Related to signal/noise (S/N) ratio.

BLOS. Beyond line of sight.
Long distance communications. Generally requires HF propagation, relay, or specialized high power, lower frequency transmission.

BX.25.
The AT&T variant of the X.25 standard.

Byte or octet.
8 bits. One character representation.

CAD. Collective Address Designation.
A task group, for instance, has a collective group call sign.

Catenet.
See Internetwork.

Codec. Coder-decoder.
A codec provides error coding data in a transmission channel and forward error corrects data in a receiving channel.

Communications Station.
In this thesis, any node that acts as a network controller is a communications station. This is the network hub or network control station. See ship.

CPU. Central Processing Unit.
For the computers discussed in this thesis, the CPU is a single chip. A microprocessor contains the CPU (which in turn contains registers, an arithmetic/logic unit, and control), RAM and I/O. A raw measure of CPU power is clock speed in mHz.

CRC. Cyclic Redundancy Checksum.
A mathematical algorithm when applied to a block of data yields a 2 or 4 byte figure that can be duplicated by a receiver. In ARQ (backward error correction) systems, receivers and transmitters both calculate CRCs on the same data. If they agree, the receiver has received the same data that the transmitter sent.

Cryptanalysis.
An interceptor decrypting and reading communications.

CSMA/CD. Carrier Sense Multiple Access/Collision Detection.
A communications scheme where several users on a channel randomly access the curcuit and determine whether their transmissions were interfered with by another transmission (which results in retransmissions). See LBT, LWT.

CUDIXS. Common User Digital Information eXchange System.
Naval ship-shore satellite communications system. CUDIXS is used as representative of several satellite systems in this thesis. Functionally, the remaining ones are the same:

- SSIXS - Submarine Satellite IXS
- ASWIXS - Anti-Submarine Warfare IXS
- TACINTEL - Tactical Intelligence IXS
- TADIX - Tactical Data IXS, major command network
- OTCIXS - Officer in Tactical Command IXS.

Datagram.
See Message.

DCE. Data Communications Equipment.
In a master/secondary communications system, the master is known as DCE. The secondaries are DTEs. See communications station.

DCS. Defense Communications Systems.
The common carrier for Department of Defense.

DDN. Defense Data Network.
Packet Switched communications network for operational DoD use. DDN grew from the experimental ARPANET and will eventually replace AUTODIN.

DES. Data Encryption Standard.
A non-classified encryption standard for cryptographic protection of communications. At present, DES can be used for data that requires protection but is not classified (Confidential, Secret, Top Secret). See KG.

DF. Direction Finding.
Part of ESM involving a passive receiver gaining a bearing on a transmitter.

DTE. Data Terminal Equipment.
In a master/secondary communications system, the secondaries are DTEs. In this thesis, I refer to DTEs as ships.

Duplex.
A communications system where two separate frequencies are used simultaneously. In this thesis, this means that the communications station sends on one frequency and the ship stations all use a second frequency. See Full Duplex and Half Duplex.

ELINT. Electronic Intelligence.
This category of ESM involves collecting and analyzing electronic emanations. An example would be fingerprinting a radar.

ELOS. Extended line of sight.
Communications beyond the horizon, but still essentially local. See LOS and BLOS.

EMCON. Emission Control.
The process of limiting ESM vulnerability by controlling own force electromagnetting emanations.

ESM. Electronic support Measures.
Intelligence collection by passively intercepting an enemy's emissions. Includes DF, SIGINT, ELINT. An ESM receiver makes no emissions of its own and can thus operate covertly.

FEC. Forward Error Correction.
The technique of adding redundancy to a bit stream in order to use the context (the added bits) to correct any errors.

Full Duplex.
A duplex system where all users are capable of sending and receiving simultaneously.

Full_ARQ. An acknowledgement procedure where each packet is acknowledged. The sender assumes that the packet has not been received until a receipt, an ACK is returned. See NAK_only and No_ack.

Go back N.
A backward error correction scheme where a batch of packets are sent. The receiver acknowledges up through the last packet received correctly. The transmitter then starts the next batch with the packet immediately after the last acknowledgement.

Half Duplex.
A duplex system where two communications channels are used but at least some users (usually the ships) cannot send and receive simultaneously.

HF. High Frequency.
Band in electromagnetic spectrum 2-30mHz where long distance communications is possible using skywave propagation (ionospheric bounce). Higher frequencies (VHF, UHF, SHF, EHF) are line of sight and require a satellite relay for long distance terrestrial propagation. Lower frequencies (MF, LF, VLF, ELF) have good groundwave characteristics (use lithosphere as a waveguide) but require antenna sizes and power outputs that make them impractical for shipboard transmitting use.

HFSS. High Frequency Ship-Shore.
A term used by Naval Research Lab to distinguish BLOS systems from LOS and ELOS systems and Intra-Task-Force (ITF) communications.

I/O. Input/Output.
That part of a computer operating system that deals with communications with peripheral devices such as terminals and communications links.

Internetwork.
A structure of interconnected local area networks. This differs from wide area networks in that local networks are simply connected to each other by bridges, protocol converters or network interface units. No backbone network emerges.

IP. Internet Protocol.
A set of procedures for interconnecting multiple LANs.

ISO. International Standards Organization.

KG. Key Generator.
A machine cypher device that uses 1) a key and 2) input data to generate output data. If the input data is plaintext, the output is codetext; if the input is codetext, the output is plaintext. At present, KGs are used for military classified communications. See DES.

LAN. Local Area Network.
Communications system where all nodes are connected directly to each other by a bus, chained together in a ring, or all directly connected in a star arrangement to a central node. Topology defines a LAN, not distance. Our HF system is a star LAN in topology.

LBT. Listen Before Talk.
A network access scheme involving listening on a channel to see if it is clear before transmitting. See CSMA/CD.

LDMX. Local Digital Message eXchange.
Major command message processing system. RIXT circuits are usually connected to an LDMX or NAVCOMPARS processor. Functionally, LDMS and NAVCOMPARS use the same equipment and software.

LOS. Line of Sight.
Frequencies above HF are generally LOS frequencies. Transmitting and receiving antennae must 'see' each other to communicate.

LPI. Low Probability of Intercept.
A factor of EMCON where a communicator offers a low opportunity for enemy ESM to intercept his communications.

LUF. Lowest Usable Frequency.
The lowest HF frequency for which ionospheric conditions will support a specific long distance channel. The LUF is somewhat dependent on the amount of power transmitted and antenna efficiency.

LWT. Listen While Talk.
A network user listens while transmitting to see if other users are interfering.

Message.
A logical unit of information. In the ARPANET community, the term datagram is used.

Modem. Modulator-demodulator.
A modem translates baseband (digital) signals into analog signals and back. (This is not the modulator within a radio that modulates an analog signal onto a carrier frequency for transmission).

MUF. Maximum Usable Frequency.
The highest HF frequency for which ionospheric conditions will propagate a signal between two radio stations. Unlike the LUF which is somewhat power dependent, the MUF is dependent on ionospheric conditions and the distance between the two stations.

Multiplex.
Two definitions of multiplex exist:

1) downward multiplexing is the process of getting multiple users operating on a single channel.

2) upward multiplexing is the process of one (or more) users using multiple channels to communicate more data than a single channel would carry.

NAK. Negative AcKnowledgement.
Message not received. Derived from ASCII character 15H, ^M. See ZDK.

NAK_only. An acknowledgement procedure where the sender assumes that a receiver received a packet unless it receives a NAK (ZDK request). See Full_ARQ and No_ack.

NAVCOMPARS. Naval Communications Processing and Routing System.
Network interface unit at a NAVCAMS that serves to interconnect fleet communications circuits with Autodin Switching Centers.

NCS. National Communications System. Naval Communications Station.
The term NECOS (Network Control Station) or simply 'communications station' (commsta) is used for naval communications station. National Communications System includes all assets available for non-military governmental and civil communications in the event of mobilization or disaster.

NECOS. Network Control Station.
See NCS and DCE and communications station.

No_ack. A state where a packet is not acknowledged by a receiver. This procedure is used when perishable, refreshed data is being distributed; it is more effective to await the next edition of the packet than to attempt to retrieve a damaged one. See Full_ARQ and NAK_only.

Octet. See Byte.

OSI. Open System Interconnection model.
A reference model for interconnection of communicating computers.

Packet.
A physical unit of information.

Q-signals.
A radio brevity code with specialized communications meanings. All of these codes start with q or z, thus the names q- and z-signals.

QSL.
A radio Q-signal indicating receipt of a message. See ACK.

RAM. Random Access Memory.
In contemporary usage, this is solid state electronic memory and is measured in bytes (one byte is 8 bits and equivalent to one character representation) and kilobytes (1000 bytes). RAM is one raw measure of CPU power of a computer.

RATS. Radio Amateur Telecommunication Society.

RATS. Random Access Time Slot.
A silent period during which no scheduled transmissions take place. A user wishing to enter a network transmits his entry request during this period on a CSMA/CD basis. The term Silent Period is used in this thesis.

RIXT. Remote Information eXchange Terminal.
A Navy local automated message switching circuit designed to provide rapid message service to several users in a local area.

RS-232. ANSI Recommended Standard number 232.
A common protocol for asynchronous digital baseband data transfer.

S/N. Signal to Noise Ratio.
An analog measure of quality and capacity of a communications channel.

SAMPS. Semi-Automated Message Processing System.
Coast Guard's local automatic communications circuit. Usually used to link shore units and inport ships in a district.

Selective repeat.
A backward error correction scheme where a batch of packets are sent. The receiver acknowledges all correct packets. The transmitter then resends unacknowledged packets and any new packets where there is room.

Ship.
In this thesis, the term ship is appled to any communicating node except the communications station. This shorthand term could mean ship, submarine, aircraft, served shore station. See Communications Station.

SIGINT. Signals Intelligence.
A branch of ESM targeting specifically on an enemy's communications. Includes cryptanalysis, traffic analysis, traffic flow analysis, and direction finding.

Silent Period.
A designated period of time where routine traffic is suspended and all stations listen for emergency traffic. The distress frequencies 500kHz (CW) and 2182kHz (voice) both have two silent periods each hour. In this thesis, silent periods are those times where no traffic is scheduled. Users wishing to enter the network use this silent period to send their network access requests.

Simplex.
A communications system where all users share the same frequency and only one can send at a time. A user cannot send and receive at the same time. See Duplex and Multiplex.

SOL. Sequence Order List.
A schedule maintained by the communications station indicating which ships can transmit when.

Stop and wait.
A backward error correction scheme where each packet is individually acknowledged before proceding.

T1. Primary Timer.
The timer associated with a packet when it is sent. If the T1 times out without the sender receiving an ACK or NAK, the packet is resent. See the X.25 standard, part 2.4.7.

TCP. Transport Control Protocol.
Essentially a communications oriented operating system for a computer that supports internetworking.

TNC. Terminal Node Controller.
The amateur radio Link Level controller that is the centerpiece of amateur packet radio.

Traffic analysis.
An intelligence process where the interceptor is not actually reading the contents of messages but gaining insight into enemy order of battle and intentions by examining who is talking to whom.

Traffic flow analysis.
An intelligence process where traffic flow is monitored. An interceptor will not know who is talking with whom or what they are saying, but inferences may be drawn from the volume of traffic on a circuit.

WAN. Wide Area Network.
Backbone trunking structure that connects multiple Local Area Networks together. Autodin in an example of a WAN.

X.25.
CCITT standard dealing with ISO levels 1,2 and 3 for computer communications.

Z-signals. See Q-signals.

ZDK.
A radio Z-signal indicating non-receipt of a message and request for retransmission. See NAK.

BIBLIOGRAPHY

[Baker-Ephremides-Wieselthier, 82] D. J. Baker, A. Ephremides, J. E. Wieselthier, "An Architecture for the High-Frequency Intratask Force (ITF) Communications Network", NRL Report 8638, Dec 1982.

[Baker-Wieselthier-Ephremides, 84] D. J. Baker, J. E. Wieselthier, A. Ephremides, D. N. McGregor, "Resistance to HF Jamming Interference in Mobile Radio Networks by an Adaptive, Distributed Reconfiguration Technique", NRL Report 8834, Aug 1984.

[Bauman et al, 78] Bauman, R. M., Golliday, C. L., Royce, R. K., Andrews, D. C., Hobbis, C. E., "Adaptive Cancellation of Local Electromagnetic Interference in Naval HF Communication Systems", NRL Report number 8175, Jul 78.

[Behre, 85] Behre, C. P., "Proposal for an Automated Intra-Task Force Narrative Record Information Exchange", Naval Postgraduate School (Masters Thesis), Mar 85.

[Berlekamp, 74] Elwyn R Berlekamp ed., The Development of Coding Theory, IEEE Press, 1974.

[Boehm, 81] Boehm, Barry W. Software Engineering Economics, Prentice-Hall, 1981.

[Brayer, 81] Kenneth Brayer, "Error Correction Code Performance on HF, Troposcatter, and Satellite Channels", IEEE Transactions on Communications Technology, October 1981, Vol Com-19 #5, p 781.

[Brooks, 75] Brooks, Frederick P. Jr. The Mythical Man-Month: Essays on Software Engineering, Addison-Wesley, 1975.

[CCITT, 84] CCITT, X.25, "Interface between Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE) for terminals operating in the packet mode and connected to public data networks by dedicated circuits.", International Telecommunications Union, Geneva 1976, amended at Geneva 1980 and Malaga-Torremolinos 1984.

[Contel, 82] "Final Report: Communications Protocol Structure", Contel Information Systems, Vienna, Va, Sep 1982.

[Fairley, 85] Software Engineering Concepts, McGraw-Hill, 1985.

[Hauser, 84] J. Hauser, McGregor & Baker, Design and Simulation of an HF Ship-Shore Communication Network Protocol, Naval Research Laboratory Report 8805, August 1984.

[Kozaczuk, 79] Kozaczuk, Wladyslaw, Enigma: How the German Machine Cipher Was Broken and How It Was Read by the Allies in World War Two, University Publications of America, Inc. 1979.

[Lin-Costello, 83] Shu Lin, Daniel J Costello, Jr, Error Control Coding: Fundamentals and Applications, Prentice-Hall, 1983.

[LORAN, 80] US Coast Guard, Loran-C User Handbook, COMDTINST M16562.3, May 1980.

[LORAN, 81] US Coast Guard Specification of the Transmitted Loran-C Signal, COMDTINST M16562.4, July 1981.

[Melich-Landwehr-Crepeau, 80] Melich, Michael E., Landwehr, Carl E., Crepeau, Paul J., "Analysis of Alternative Satellite Channel Management Systems", NRL Report # 8404, Oct 1980.

[NavMat, 85] Naval Material Command, Navy Program Manager's Guide, 1985 Edition, P-9494, Washington DC 20360.

[Nichols, 82] Elizabeth Nichols, Joseph Nichols & Keith Musson, Data Communications for Microcomputers, McGraw-Hill, 1982.

[Price, 85a] Price, Harold, "What's all this Racket about Packet?", Jul 85, p 14 and "A Closer Look at Packet Radio", Aug 85, p 17, QST.

[Price, 85b] Price, Harold NK6K, "What's All this Racket about Packet", Jul 85, p 14. and "A Closer Look at Packet Radio", Aug 85, p. 17. QST.

[RFC___.] Documentation for protocols as well as issues relating to ARPANET development are maintained on-line by Stanford Research Institute in a library of Requests For Comments. The list is updated periodically. The two RFCs used in this thesis are:

RFC791 -- Internet Protocol Specification

RFC793 -- Transmission Control Protocol Specification

Both documents were prepared for DARPA by Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina del Rey, Ca 90291, September 1981.

[Rosner, 82] Roy D. Rosner, Packet Switching: Tomorrow's Communications Today, Wadsworth, 1982.

[Stallings, 84] William Stallings, Local Networks: An Introduction, MacMillan, 1984.

[Stallings, 85] William Stallings, Data and Computer Communications, MacMillan, 1985.

[Stremler, 82] Ferrel G Stremler, Introduction to Communication Systems, Addison-Wesley, 1982.

[Tanenbaum, 81] Andrew S. Tanenbaum, Computer Networks, Prentice-Hall, 1981.

[Tobagi, 84] Fouad Tobagi, Richard Binder, Barry Leiner, "Packet Radio and Satellite Networks", IEEE Communications Magazine, November 1984, p 24.

[Wakerly, 78] John Wakerly, Error Detecting Codes, Self-Checking Circuits and Applications, Elsevier North-Holland Inc, 1978.

[Wieselthier-Baker-Ephremides, 81] J. E. Wieselthier, D. J. Baker, A. Ephremides, "Survey of Problems in the Design of an HF Intra Task Force Communication Network", NRL Report 8501, Oct 1981.

INITIAL DISTRIBUTION LIST

		Copies
1.	Defense Technical Information Center Cameron Station Alexandria, Va 22304-6145	2
2.	Library, Code 0142 Naval Postgraduate School Monterey, Ca 93943-5000	2
3.	Lcdr Rex A Buddenberg Commandant (G-TES-1) US Coast Guard Washington, DC 20593	2
4.	Commanding Officer Coast Guard Station 7323 Telegraph Road Alexandria, Va 22310	1
5.	Commanding Officer US Coast Guard Electronics Engineering Center Wildwood, NJ 08260	1
6.	Commanding Officer Coast Guard Communications Station San Francisco - NMC PO Box 560 Point Reyes Station, Ca 94956	1
7.	Commander Space and Naval Warfare Systems Command Code 155-1 HF Tactical Systems Washington, DC 20363-5100	1
8.	Naval Telecommunications Systems Integration Center Cheltenham Washington, DC 20390	2
9.	Commander Third Fleet Pearl Harbor, Hi 96860 Attn: Cdr Berthiaume	1
10.	Commander in Chief Pacific Fleet Makalapa, Hi 96860	1
11.	Commander Naval Ocean Systems Center San Diego, Ca 92152-5000 Attn: Code 772, Bob Rose	1

12. Commander
Naval Ocean Systems Center
San Diego, Ca 92152-5000
Attn: Code 84, K. R. Casey
1

13. Naval Research Laboratory
Ray Cole
Code 7523
Washington, DC 20375
1

14. Mr. Ken Boheim
NCS/PP
Eighth Street and South Courthouse Road
Arlington, Va 22204
1

15. Mr. Edward H. Cain
NCS/PP
Eighth Street and South Courthouse Road
Arlington, Va 22204
1

16. Dr. Bruce Barrow
NCS/PP
Eighth Street and South Courthouse Road
Arlington, Va 22204
1

17. Col. William Schooler
NCS/EP
Eighth Street and South Courthouse Road
Arlington, Va 22204
1

18. Mr. Norman Douglas
NCS/EP
Eighth Street and South Courthouse Road
Arlington, Va 22204
1

19. LtC Tom Cindric (JDSSC)
Defense Communications Agency
Code 662
Washington, DC 20305
1

20. Dr. Jack W. LaPatra
Telemedia Inc.
310 South Michigan Ave.
Chicago, Il 60604
1

21. MITRE Corporation
1820 Dolley Madison Boulevard
McLean, Va 22102
Attn: Fred Leiner
1

22. Cpt Brad Bryant 1
DCA
C45 Code A550
Building A
Arlington Hall Station
Arlington, Va 22212-5410

23. Director 1
National Security Agency
9800 Savage Road
Ft Meade, Md 20755-6000
Attn: V311